

Full Circle Magazine



Issue #153

Table of Contents

1. [Full Circle Magazine](#)
2. [Editorial](#)
3. [News](#)
4. [Command & Conquer](#)
5. [How To - Python in the REAL World - Part 101](#)
6. [How-To - Test Linux in VirtualBox](#)
7. [How-To - Darktable - Part 11](#)
8. [The Daily Waddle](#)
9. [How To - Drawing with Inkscape - Part 93](#)
10. [Using Krita to Rework Old Photos - Part 3](#)
11. [The Daily Waddle](#)
12. [Linux Loopback](#)
13. [Everyday Ubuntu - Personal Information Management Part 1](#)
14. [Ubports Devices](#)
15. [The Daily Waddle](#)
16. [My Opinion - Container Security Concerns](#)
17. [My Opinion - Linux on the Go](#)
18. [How To - Write For Full Circle Magazine](#)
19. [Review - mtPaint Graphics Editor](#)
20. [Interview: Jaron Parsons - Fury BSD Dev](#)
21. [Letters](#)
22. [Q&A](#)
23. [Linux Certified - Part 9](#)
24. [Ubuntu Games - Stygian – Reign of the old ones](#)
25. [Patreon](#)
26. [Patrons and Donors](#)
27. [How to Contribute](#)

Landmarks

1. [Table of Contents](#)



The articles contained in this magazine are released under the Creative Commons Attribution-Share Alike 3.0 Unported license. This means you can adapt, copy, distribute and transmit the articles but only under the following conditions: you must attribute the work to the original author in some way (at least a name, email or URL) and to this magazine by name ('Full Circle Magazine') and the URL www.fullcirclemagazine.org (but not attribute the article(s) in any way that suggests that they endorse you or your use of the work). If you alter, transform, or build upon this work, you must distribute the resulting work under the same, similar or a compatible license. Full Circle magazine is entirely independent of Canonical, the sponsor of the Ubuntu projects, and the views and opinions in the magazine should in no way be assumed to have Canonical endorsement.

Please note: articles in this magazine are provided with absolutely no warranty whatsoever; neither the contributors nor Full Circle Magazine accept any responsibility or liability for loss or damage resulting from readers choosing to apply this content to theirs or others computers and equipment.

RETURN TO CONTENTS

Editorial

Welcome to the latest issue of Full Circle!

Welcome to the first issue of 2020! We bring in the roaring twenties with the usual suspects of Python, Darktable, Inkscape, and Krita. And if your new year's resolution is to try new distros then you'll love the HowTo on testing Linux in VirtualBox.

SJ brings us a BSD combo with a quick look at Project Trident and other BSD options and then an interview with one of the FuryBSD developers.

We have a couple of opinion pieces this month too. One is on container security concerns, and the other is about Linux on the go. It's not for the faint of heart as it mentions the M-word.

I won't keep you any longer. Read on!

All the best to you and yours for 2020!

Ronnie

ronnie@fullcirclemagazine.org



FCM PATREON : <https://www.patreon.com/fullcirclemagazine>



[RETURN TO CONTENTS](#)

NEWS

Submitted by Arnfried Walbrecht

Full Circle Weekly News



<http://fullcirclemagazine.org/feed/podcast>

[http://www.stitcher.com/s?
fid=85347&refid=stpr](http://www.stitcher.com/s?fid=85347&refid=stpr)

[http://tunein.com/radio/Full-Circle-
Weekly-News-p855064/](http://tunein.com/radio/Full-Circle-Weekly-News-p855064/)

[https://player.fm/series/the-full-circle-
weekly-news](https://player.fm/series/the-full-circle-weekly-news)

Reiser5 file system available:

Reiser5 doesn't implement its own block layer like ZFS etc. In the reiser approach scaling out is performed by file system means, rather than by block layer means. The flow of IO-requests issued against each device is controlled by the user. To add a device to a logical volume with parallel scaling out, you first need to format that device - this is the difference between parallel and non-parallel scaling at first glance.

Systems with parallel scaling out provide better scalability and resolve a number of problems inherent to non-parallel ones, like RAID "bottlenecks".

The whole thing is explained in detail in the newsletter.

<https://marc.info/?l=reiserfs-devel&m=157780043509663&w=2>

Bonsai, a device synchronization service for GNOME:

Christian Hergert, the author of the integrated environment GNOME Builder, now at Red Hat, introduced his pilot project Bonsai, aimed at solving the problem of: synchronizing contents of several devices that use GNOME. Users can use Bonsai to link several Linux devices on their home network, when they need to access files and application data on all computers, but do not want to transfer their data to third-party cloud services. The project code is written in C and is released under the GPLv3 license.

<https://blogs.gnome.org/cherbert/2020/01/01/introducing-bonsai/>

GitBucket 4.33 Collaborative Development System Released:

GitBucket 4.33 project has been released, a framework for working with the Git repository, which provides an interface in the style of GitHub and Bitbucket. The system is notable for

easy installation, the ability to expand functionality through plugins and compatibility with the GitHub API. The code is written in Scala and is licensed under Apache 2.0. MySQL and PostgreSQL can be used as DBMS.

<https://gitbucket.github.io/gitbucket-news/gitbucket/2019/12/31/gitbucket-4.33.0.html>

Final release of Supertux Kart 1.1:

Release SuperTuxKart 1.1, a free racing game with plenty of go-karts, tracks and features. The game code is distributed under the GPLv3 license. Binary assemblies are available for Linux, Android, Windows and macOS.

<http://blog.supertuxkart.net/2020/01/supertuxkart-11-released.html>

Arch Linux moved to zstd for packet compression:

Arch Linux developers have announced that they are moving the packaging scheme from the xz algorithm (.pkg.tar.xz) to zstd (.pkg.tar.zst). The reassembly of packages in zstd format led to a total increase in packet size by 0.8%, but sped up unpacking time by 1300%. As a result, switching to zstd will lead to a noticeable increase in package installation speed. Currently, 545 packages are already compressed in the repository using the zstd algorithm, the remaining packages will be transferred to zstd as updates are generated for them.

Packages in the .pkg.tar.zst format are collected automatically when using devtools 20191227 and newer releases of this toolkit. For users, the transition to a new format does not require intervention, pacman (5.2) and libarchive (3.3.3-1, released back in 2018) were already updated last year . For those with a non-updated release of libarchive, the new version can be installed from a separate repository .

<https://www.archlinux.org/news/now-using-zstandard-instead-of-xz-for-package-compression/>

Portable release of CWM 6.6 window manager:

A portable version, 6.6 of the CWM window manager, developed as part of the OpenBSD project, was released . This window manager is based on the evilwm code , but uses modern X11 protocol interfaces, and, traditionally for OpenBSD, it has been developed with special attention to security issues. The portable release, in addition to OpenBSD, supports the operating systems FreeBSD, NetBSD, macOS (version 10.9 and higher), as well as operating systems based on the Linux kernel.

<https://github.com/leahneukirchen/cwm>

Smartmontools 7.1:

A new version of the smartmontools 7.1 package has been released, containing smartctl and smartd applications for monitoring and controlling (S)ATA, SCSI / SAS and NVMe drives that support SMART technology. It works on Linux, FreeBSD, Darwin (macOS), Windows, QNX, OS / 2, Solaris, NetBSD, and OpenBSD.

<https://www.smartmontools.org>

Fedora 32 intends to enable earlyoom:

A plan has been published to include the default earlyoom background process in Fedora 32 to respond early to a lack of memory on the system. If the amount of available memory is less than the specified value, then earlyoom through sending SIGTERM (free memory less than 10%) or SIGKILL (< 5%) forcibly (terminates the process that consumes the most memory (having the highest value / proc / * / oom_score), without bringing the state of the system to clearing the system buffers, Earlyoom will allow you to respond more quickly to a lack of memory without having to call the OOM (Out Of Memory) handler in the kernel, which is triggered when the situation becomes critical and the system, as a rule, no longer responds to user actions.

<https://fedoraproject.org/wiki/Changes/EnableEarlyoom>

Void Linux project trident beta:

The first beta version of the Trident OS is available, migrated from FreeBSD and TrueOS to the Void Linux package base. The size of the bootable iso image is 515MB. The distro uses ZFS on the root partition and it is possible to roll back the boot environment using ZFS snapshots. A simplified installer is supplied, that can work on systems with EFI and BIOS. Encryption of your swap partition is possible. Package options for standard glibc and musl libraries are offered, each user gets a Separate ZFS dataset for the home directory (you can manipulate snapshots of the home directory without gaining root privileges), data encryption in user directories is also provided.

<https://project-trident.org/post/void-beta-available/>

Vim9, fork of Vim for script optimization experiments:

Bram Moolenaar, author of Vim, announced the creation of a repository for Vim9. It is an experimental fork of Vim, aimed at

exploring possible ways to improve the productivity and quality of work of Vim scripting language.

The main optimizations are related to the processing of methods for determining, calling and executing functions, as well as avoiding the use of dictionaries for arguments and local variables. The initial prototype of the new implementation, in which functions are first compiled into a sequence of instructions, that store intermediate results and local variables on the stack, reduced the time it takes to perform a cyclic function call test from 5.018541 to 0.073595 seconds, and a string processing test from 0.853752 to 0.190276 seconds. Vim9 also develops tools for writing plugins, not only in the built-in scripting language, but also in various programming languages, including Python, Go, and Java.

https://groups.google.com/forum/#!msg/vim_dev/OPbZwpcBP98/n4A1cviUBwAJ

Gdevelop 5.0 beta released:

GDevelop is a full-featured, open-source game development software, allowing to create HTML5 and native games without any knowledge in a specific programming language. All the game logic is built up using an intuitive and powerful event-based system.

GDevelop 5.0.0-beta84 adds in a bunch of new effects you can use including “Black and White, Noise, CRT, Godray, Tilt shift, Advanced bloom, Kawase blur, Zoom blur, Displacement, Color Map” and more. These special effects can be added to your game with minimal fuss too, and the result is pretty awesome.

<https://gdevelop-app.com/>

Tablet Friendly Program MyPaint Finally Hits Beta:

MyPaint is a drawing, painting program dedicated to pressure-sensitive devices, tablets. This program is a must for digital painters which offers distraction-free drawing environments with tons of new features.

This free and open-source program packs important features such as versatile and configurable brush engine, customer brush creation and also support for Wacom and other devices. After almost 10 years since version 1.0, MyPaint 2.0 beta is released. This beta release is an important milestone for the upcoming release.

<http://mypaint.org/>

Staying with paint, VPaint 1.7 released:

VPaint is an experimental prototype based on the Vector Graphics Complex (VGC), a technology developed by a collaboration of researchers at Inria and the University of British Columbia. It allows you to create resolution-independent illustrations and animations using innovative techniques.

After four years of development, VPaint 1.7 has been released, combining a vector graphics editor and a system for creating 2D animation. The program is positioned as a research project with an experimental implementation of the mathematical concept of VGC (Vector Graphics Complex), which allows you to create animations and illustrations that are not tied to pixel resolution. The project's achievements are written in C++ (using the Qt and GLU libraries) and are distributed under the Apache 2.0 license. Builds are available for Linux (AppImage), Windows and macOS.

In the future, VPaint aims to compete with Adobe Illustrator, Autodesk Graphic, CorelDRAW and Inkscape, and then secondary with Adobe Animate, ToonBoom Harmony, CACANi, Synfig and OpenToonz. Both packages, despite having a price tag, will ship as open source under the Apache 2.0 license. Builds for Linux will ship for free (only editions for Windows and macOS will be paid).

<http://www.vpaint.org>

Valve fixes an error on Steam that led to an incorrect calculation of the share of Linux users:

A bug was fixed in which the Linux build of the program which crashed while preparing data about the user's environment, This is where the Steam Hardware & Software Survey statistics are calculated. It is assumed that the resolved problem was one of the reasons for the low share of Linux in Steam reports, since the attempt to send data from Linux users crashed. In the December Steam report, Linux accounted for just 0.67%.

<https://steamcommunity.com/groups/steamworks/announcements/detail/1591381408652851752>

VVVVVV game open sourced:

Terry Cavanagh posted the VVVVVV source code after a decade of playing. VVVVVV is a platform game with graphics in the style of the old Atari 2600, where the player can change the direction of gravity instead of jumping.

The source code of two versions of the game are available - for desktop systems in C + + and for mobile platforms on Actionscript for Adobe AIR. The code is hosted on GitHub under its own limited license, VVVVVV Source Code License v1.0 , which does not allow you to modify or distribute the code for commercial purposes. Game resources, graphics and music remains under a proprietary license that allows free use for personal purposes.

<https://github.com/TerryCavanagh/VVVVVV/blob/master/LICENSE.md>

PHP framework: Laravel 6.10 - bug fixes and new features:

PHP framework Laravel, version 6.10, with an MIT license, has been released. The developers report on eleven changes. Laravel is also said to support PHPUnit v9, which is expected in February.

The latter is also justified by the fact that Laravel 6.x is a version with long-term support.

The makers include the macro “validateWithBag”, which enables the error bag to be specified when validating a request. The macro can also be found in the documentation, writes developer Paul Redmond. Macros for the Redis Connection have also been added.

In the bug fixes, the developers list repaired code for `getenv ()`. The validation rules for `exclude_if \ exclude_unless` have also been repaired. The release notes provide further details.

<https://github.com/laravel/framework/blob/6b9232037f8041c298b6479ef4ffd70d78d61a27/CHANGELOG-6.x.md#v6100-2020-01-07>

Endless OS 3.7.5: offline distribution release:

Formerly specially designed for emerging markets, Endless OS’s content is primarily aimed at families and can also be used offline to a large extent. With an easy-to-use desktop, pre-installed games and educational software, the project’s creators want to win over beginners in particular for their operating system.

<https://endlessos.com/computers/>

MariaDB: Ex-Tesla manager bolsters the management team:

MariaDB Corporation announces that Susan Repo is now Chief Operational Officer (COO). Repo can look back on a long and successful career, which includes leading an auto-fintech start-up and five years in top positions at Tesla. There Repo steered the manufacturer of electric cars through its exorbitant growth phase. At MariaDB, Repo will bring her extensive experience to monitor and accelerate business operations.

Susan Repo most recently held the positions of CEO and President of DriveOn, an auto fintech startup. Previously, she served as CFO and Vice President of Finance at Tesla, where she raised over \$ 11 billion in capital and contributed to key legal and governance initiatives within Tesla’s global

organization. Repo expanded the financial sector and played a leading role in the acquisition of SolarCity in 2016. Prior to this, Repo held various senior legal and finance roles in the technology industry, including Juniper Networks and Agilent Technologies. All-in-all, it seems like a good appointment.

<https://mariadb.com/>

Version 1.0 for Hummingbird VPN client:

VPN client Hummingbird from AirVPN is ready for production with version 1.0. The client runs on Linux and Raspi and is based on a fork from OpenVPN 3.

Hummingbird 1.0 is available at Gitlab, there are versions for Linux, the Raspberry Pi and Mac OS, but none for Windows. AirVPN offers the software as a standalone client under GPLv3, it should be particularly fast and at the same time only leave a small footprint in the system's working memory.

Hummingbird 1.0, which in addition to numerous bug fixes also offers some additional features. The CHACHA20-POLY1305 cipher is optionally used for both control and data channels, which speeds up the performance on ARM and Linux-based platforms that do not support AES-NI. The cipher can be selected via “-cipher name” when the client is started.

<https://gitlab.com/AirVPN/hummingbird/tree/master>

Ubuntu GamePack 18.04 released:

The Ubuntu GamePack 18.04 build is available for download , which includes tools for launching more than 55 thousand games. Applications include those specially designed for the GNU / Linux platform, and Windows games launched using PlayOnLinux, CrossOver and Wine, as well as old games for MS-DOS. The distribution is based on Ubuntu 18.04 and includes all updates as of January 2020. By default, the GNOME Flashback interface is the default DE, closely resembling the classic GNOME shell, but other environments are available. The size of the iso image is 4.1 GB (x86_64). An

update is also available for the previous branch based on Ubuntu 16.04, which was also compiled for 32-bit i386 systems.

<https://ualinux.com/ru/ubuntu-gamepack>

Solaris 11.4 SRU 17 Update:

The Solaris 11.4 SRU 17 (Support Repository Update) operating system update has been published , which offers a series of regular fixes and improvements for the Solaris 11.4 branch . To install the fixes proposed in the update, just run the 'pkg update' command.

link has been shagged by oracle login

CentOS 8.1 (1911) release:

The CentOS 1911 distribution release introduces changes from Red Hat Enterprise Linux 8.1. The distribution is fully binary compatible with RHEL 8.1, changes made to packages are usually reduced to rebranding and replacement of decorations. Builds of CentOS 1911 are available (7 GB DVD and 550 MB netboot) for x86_64, Aarch64 (ARM64) and ppc64le architectures. SRPMS packages, the way the binary files were assembled, and debuginfo are available through vault.centos.org.

<https://wiki.centos.org/Manuals/ReleaseNotes/CentOS8.1911>

Vulkan 1.2:

The graphics standard development consortium, Khronos, has published the Vulkan 1.2 specification, which defines an API for accessing the graphics and computing capabilities of your GPU. The new specification has incorporated fixes and extensions accumulated over two years. Drivers supporting the new Vulkan version have already been released by Intel, AMD, ARM, Imagination Technologies and NVIDIA . In Mesa, Vulkan 1.2 support is offered for RADV drivers (AMD cards) and ANV (Intel) drivers . Vulkan 1.2 support is also implemented in the

RenderDoc 1.6 debugger, the LunarG Vulkan SDK, and a sample setVulkan-Samples already.

<https://www.khronos.org/news/press/khronos-group-releases-vulkan-1.2>

First stable release of Fedora CoreOS:

Fedora project developers have announced the stabilization of the Fedora CoreOS distribution and its availability for widespread use. Fedora CoreOS is being promoted as a single solution for running environments based on isolated containers, replacing Fedora Atomic Host and CoreOS Container Linux products. Support for the CoreOS Container Linux distribution will be discontinued after 6 months, and support for Fedora Atomic Host is expected to end in late November.

Fedora CoreOS aims to provide a minimal environment, atomically updated automatically without the participation of an administrator and unified for the mass deployment of server systems designed exclusively to run containers. The distribution package provides only the minimum set of components sufficient to run isolated containers - the Linux kernel, the systemd system manager and a set of service services for connecting via SSH, configuration management, and installing updates.

<https://fedoramagazine.org/fedora-coreos-out-of-preview/>

PinePhone Linux smartphone is available for order:

Pine announced the start of deliveries to all interested parties. The first smartphone Limited PinePhone (Braveheart Edition), developed by the community Pine64 (note: the first batch is already sold out). The start of mass production is scheduled for March 2020. As stated originally, the cost of a smartphone is \$ 150. The device is designed for enthusiasts who are tired of Android and want to get a fully controlled (owned?) and protected environment based on alternative open Linux platforms.

<https://www.pine64.org/2020/01/15/pinephones-start-shipping-all-you-want-to-know/>

Mir 1.7 Display Server Release:

The Mir 1.7 display server is released , the development, which is continued by Canonical, despite dropping the development of Unity shell and the Ubuntu edition for smartphones. Mir remains in demand in Canonical projects and is now positioned as a solution for embedded devices and the Internet of things (IoT). Mir can be used as a composite server for Wayland, which allows you to run any applications using Wayland in Mir-based environments (for example, those built with GTK3 / 4, Qt5 or SDL2). Installation packages are available for Ubuntu 04/16/19/20 (PPA) and Fedora 29/30/31 . The project code is distributed under the GPLv2 license.

The new release mainly offers bug fixes related to experimental support for launching X11 applications in Wayland-based environments (using Xwayland). For X11, the ability to decorate windows was implemented and an option was added to redefine the path to the Xwayland executable file. They cleaned up Xwayland related code. In one of the next issues X11 support will remove the status of the experimental function.

<https://discourse.ubuntu.com/t/mir-1-7-0-release/14048>

VirtualBox 6.1.2, 6.0.16, and 5.2.36 releases:

Maintenance releases, focussing on bug fixes. Well worth your time upgrading.

Now supports Kernel 5.5.

Swatted 18 vulnerabilities.

<https://www.mail-archive.com/vbox-announce@virtualbox.org/msg00201.html>

LXD 3.19 has been released:

This release, includes huge feature they've been working on for

the past few months, virtual machine support! It's now possible to run LXD on a system and manage both containers and virtual machines through the exact same CLI, API or even as part of a cluster deployment!

Other features include, user experience improvements.

<https://discuss.linuxcontainers.org/t/lxd-3-19-has-been-released/6529>

Another developer deletes his repository:

This time it is the author of the actix-web framework, written in Rust, who deleted the repository after being criticized for “misuse” of the Rust language. The actix-web framework, which was downloaded more than 800 thousand times, allows embedding the http-server and client functionality in Rust applications. Developing actix-web is aimed at achieving maximum performance and is the leader in many tests of web-frameworks.

<https://twitter.com/fafhrd91/status/1218135374339301378a>

PHP framework: Laravel 6.10 - bug fixes and new features:

PHP framework Laravel, version 6.10, with an MIT license, has been released. The developers report on eleven changes. Laravel is also said to support PHPUnit v9, which is expected in February.

The latter is also justified by the fact that Laravel 6.x is a version with long-term support.

The makers include the macro “validateWithBag”, which enables the error bag to be specified when validating a request. The macro can also be found in the documentation, writes developer Paul Redmond. Macros for the Redis Connection have also been added.

In the bug fixes, the developers list repaired code for `getenv ()`. The validation rules for `exclude_if \ exclude_unless` have also been repaired. The release notes provide further details.

<https://github.com/laravel/framework/blob/6b9232037f8041c298b6479ef4ffd70d78d61a27/CHANGELOG-6.x.md#v6100-2020-01-07>

Endless OS 3.7.5: offline distribution release:

Formerly specially designed for emerging markets, Endless OS’s content is primarily aimed at families and can also be used offline to a large extent. With an easy-to-use desktop, pre-installed games and educational software, the project’s creators want to win over beginners in particular for their operating system.

<https://endlessos.com/computers/>

MariaDB: Ex-Tesla manager bolsters the management team:

MariaDB Corporation announces that Susan Repo is now Chief Operational Officer (COO). Repo can look back on a long and successful career, which includes leading an auto-fintech start-up and five years in top positions at Tesla. There Repo steered the manufacturer of electric cars through its exorbitant growth phase. At MariaDB, Repo will bring her extensive experience to monitor and accelerate business operations.

Susan Repo most recently held the positions of CEO and President of DriveOn, an auto fintech startup. Previously, she served as CFO and Vice President of Finance at Tesla, where she raised over \$ 11 billion in capital and contributed to key legal and governance initiatives within Tesla's global organization. Repo expanded the financial sector and played a leading role in the acquisition of SolarCity in 2016. Prior to this, Repo held various senior legal and finance roles in the technology industry, including Juniper Networks and Agilent Technologies. All-in-all, it seems like a good appointment.

<https://mariadb.com/>

Version 1.0 for Hummingbird VPN client:

VPN client Hummingbird from AirVPN is ready for production with version 1.0. The client runs on Linux and Raspi and is based on a fork from OpenVPN 3.

Hummingbird 1.0 is available at Gitlab, there are versions for Linux, the Raspberry Pi and Mac OS, but none for Windows. AirVPN offers the software as a standalone client under GPLv3, it should be particularly fast and at the same time only leave a small footprint in the system's working memory.

Hummingbird 1.0, which in addition to numerous bug fixes also offers some additional features. The CHACHA20-POLY1305 cipher is optionally used for both control and data channels, which speeds up the performance on ARM and Linux-based platforms that do not support AES-NI. The cipher can be selected via “-cipher name” when the client is started.

<https://gitlab.com/AirVPN/hummingbird/tree/master>

Ubuntu GamePack 18.04 released:

The Ubuntu GamePack 18.04 build is available for download, which includes tools for launching more than 55 thousand games. Applications include those specially designed for the GNU / Linux platform, and Windows games launched using PlayOnLinux, CrossOver and Wine, as well as old games for MS-DOS. The distribution is based on Ubuntu 18.04 and includes all updates as of January 2020. By default, the GNOME Flashback interface is the default DE, closely resembling the classic GNOME shell, but other environments are available. The size of the iso image is 4.1 GB (x86_64). An update is also available for the previous branch based on Ubuntu 16.04, which was also compiled for 32-bit i386 systems.

<https://ualinux.com/ru/ubuntu-gamepack>

Solaris 11.4 SRU 17 Update:

The Solaris 11.4 SRU 17 (Support Repository Update) operating system update has been published , which offers a series of regular fixes and improvements for the Solaris 11.4 branch . To install the fixes proposed in the update, just run the 'pkg update' command.

link has been shagged by oracle login

CentOS 8.1 (1911) release:

The CentOS 1911 distribution release introduces changes from Red Hat Enterprise Linux 8.1. The distribution is fully binary compatible with RHEL 8.1, changes made to packages are usually reduced to rebranding and replacement of decorations. Builds of CentOS 1911 are available (7 GB DVD and 550 MB netboot) for x86_64, Aarch64 (ARM64) and ppc64le architectures. SRPMS packages, the way the binary files were assembled, and debuginfo are available through vault.centos.org

.

<https://wiki.centos.org/Manuals/ReleaseNotes/CentOS8.1911>

Vulkan 1.2:

The graphics standard development consortium, Khronos, has published the Vulkan 1.2 specification, which defines an API for accessing the graphics and computing capabilities of your GPU. The new specification has incorporated fixes and extensions accumulated over two years. Drivers supporting the new Vulkan version have already been released by Intel, AMD, ARM, Imagination Technologies and NVIDIA . In Mesa, Vulkan 1.2 support is offered for RADV drivers (AMD cards) and ANV (Intel) drivers . Vulkan 1.2 support is also implemented in the RenderDoc 1.6 debugger, the LunarG Vulkan SDK, and a sample setVulkan-Samples already.

<https://www.khronos.org/news/press/khronos-group-releases-vulkan-1.2>

First stable release of Fedora CoreOS:

Fedora project developers have announced the stabilization of the Fedora CoreOS distribution and its availability for widespread use. Fedora CoreOS is being promoted as a single solution for running environments based on isolated containers, replacing Fedora Atomic Host and CoreOS Container Linux products. Support for the CoreOS Container Linux distribution will be discontinued after 6 months, and support for Fedora Atomic Host is expected to end in late November.

Fedora CoreOS aims to provide a minimal environment, atomically updated automatically without the participation of an administrator and unified for the mass deployment of server systems designed exclusively to run containers. The distribution package provides only the minimum set of components sufficient to run isolated containers - the Linux kernel, the systemd system manager and a set of service services for connecting via SSH, configuration management, and installing updates.

<https://fedoramagazine.org/fedora-coreos-out-of-preview/>

PinePhone Linux smartphone is available for order:

Pine announced the start of deliveries to all interested parties. The first smartphone Limited PinePhone (Braveheart Edition), developed by the community Pine64 (note: the first batch is already sold out). The start of mass production is scheduled for March 2020. As stated originally, the cost of a smartphone is \$ 150. The device is designed for enthusiasts who are tired of Android and want to get a fully controlled (owned?) and protected environment based on alternative open Linux platforms.

<https://www.pine64.org/2020/01/15/pinephones-start-shipping-all-you-want-to-know/>

Mir 1.7 Display Server Release:

The Mir 1.7 display server is released , the development, which is continued by Canonical, despite dropping the development of Unity shell and the Ubuntu edition for smartphones. Mir remains in demand in Canonical projects and is now positioned as a solution for embedded devices and the Internet of things (IoT). Mir can be used as a composite server for Wayland, which allows you to run any applications using Wayland in Mir-based environments (for example, those built with GTK3 / 4, Qt5 or SDL2). Installation packages are available for Ubuntu 04/16/19/20 (PPA) and Fedora 29/30/31 . The project code is distributed under the GPLv2 license.

The new release mainly offers bug fixes related to experimental support for launching X11 applications in Wayland-based environments (using Xwayland). For X11, the ability to decorate windows was implemented and an option was added to redefine the path to the Xwayland executable file. They cleaned up Xwayland related code. In one of the next issues X11 support will remove the status of the experimental function.

<https://discourse.ubuntu.com/t/mir-1-7-0-release/14048>

VirtualBox 6.1.2, 6.0.16, and 5.2.36 releases:

Maintenance releases, focussing on bug fixes. Well worth your time upgrading. Now supports Kernel 5.5. Swatted 18 vulnerabilities.

<https://www.mail-archive.com/vbox-announce@virtualbox.org/msg00201.html>

LXD 3.19 has been released:

This release, includes huge feature they've been working on for the past few months, virtual machine support! It's now possible to run LXD on a system and manage both containers and virtual machines through the exact same CLI, API or even as part of a cluster deployment!

Other features include, user experience improvements.

<https://discuss.linuxcontainers.org/t/lxd-3-19-has-been-released/6529>

Another developer deletes his repository:

This time it is the author of the actix-web framework, written in Rust, who deleted the repository after being criticized for “misuse” of the Rust language. The actix-web framework, which was downloaded more than 800 thousand times, allows embedding the http-server and client functionality in Rust applications. Developing actix-web is aimed at achieving maximum performance and is the leader in many tests of web-frameworks.

<https://twitter.com/fafhrd91/status/1218135374339301378a>

Midnight Commander 4.8.24 file manager release:

Twin pane terminal file manger release

The most interesting improvements that landed in this release are the new file view and edit history and a fully functional

subshell in standalone editor and viewer. Apart from that a number of bugfixes and improvements have taken place touching various components such as the build system, syntax definitions and so on. Everybody is recommended to upgrade to immediately benefit from these improvements!

<https://mail.gnome.org/archives/mc-devel/2020-January/msg00008.html>

LTS versions of Qt will be available only under commercial license:

Qt Company has announced a change in its licensing model for the Qt framework, which could have a significant impact on communities and distributions using Qt. Starting with version 5.15, the QTS LTS branches will be supported until the next significant release, i.e. about half a year (updates for LTS branches are released for three years). They are hoping that such a step will accelerate the introduction of new versions and will increase the number of companies using a commercial license for Qt, the cost of which is \$5508 per year per developer (for startups and small enterprises - \$ 499 per year).

Unfortunately, the Qt policy changes are not limited to changing the license, and starting from February, you will need to register an account in the Qt Account service to download Qt binary packages. This is explained as to simplify the distribution of packages and to ensure integration with the Qt Marketplace . Access to the Jira bug tracking system, review interface, and forums will also require a Qt Account. The development and project management models remains the same .

<https://www.qt.io/blog/qt-offering-changes-2020>

Open Source Foundation publishes a petition calling for the release of Windows 7 code:

With the end of support for Windows 7 on January 14, the Free Software Foundation petitioned Microsoft's calling for transfer of Windows 7 in the free software category, to allow the community to learn and improve this OS. Microsoft has already transferred some of its programs to the open source community, but nothing significant.

This will provide evidence that Microsoft truly respects users and their freedom, instead of flowery words and marketing materials.

<https://www.fsf.org/windows/upcycle-windows-7>

Solus 4.1 released:

The release of the Solus 4.1 Linux distribution , which is not based on packages of other distributions and develops its own Budgie desktop , installer, package manager and configurator, was announced. The source code is distributed under the GPLv2 license; the C and Vala languages are used for development. Additionally, packages with GNOME, KDE Plasma, and MATE desktops are provided. The size of the iso-images is about 1.7 GB.

<https://getsol.us/2020/01/25/solus-4-1-released/>

Weston 8.0 Composite Server Release:

A stable release of the Weston 8.0 composite server has been published. This develops technologies that contribute to the full support of the Wayland protocol in Enlightenment, GNOME, KDE, and other user environments. Weston's development focuses on providing a high-quality code base and working examples for using Wayland in desktop environments and embedded solutions, such as platforms for automotive infotainment systems, smartphones, televisions, and other consumer devices. The release is scheduled for February 11.

<https://lists.freedesktop.org/archives/wayland-devel/2020-January/041147.html>

Geneva project develops automated censorship evasion:

Researchers from the University of Maryland, as part of the Geneva project attempted to create an algorithm to automate the identification of censorship methods. Manually trying to sort out the possible gaps in deep packet inspection systems (DPIs) is a rather difficult and lengthy process. Geneva tried to use a genetic algorithm to evaluate DPI features, identify errors in its implementation, and develop an optimal client-side blocking bypass strategy. The project code is written in Python.

The work of Geneva has been successfully tested to circumvent the censorship methods used in China, India and Kazakhstan. With the help of Geneva several new gaps were revealed that were not known before. At the same time, Geneva is effective only for bypassing DPI-based blocking; when blocking by IP address, it is useless and you need a VPN. During the experiments, several dozen typical DPI bypass strategies were identified that can be tested immediately without full analysis.

<https://geneva.cs.umd.edu>

SystemE, a satirical replacement for systemd with Emacs Lisp:

The Kiss Linux distribution has published the code for a satirical project systemE , positioned as a replacement for systemd, written in Emacs Lisp. The systemE toolkit allows you to organize the download using sinit as the PID 1 handler, which launches the Emacs editor under “-script” under PID2, which, in turn, executes system initialization scripts (rc.boot) written in Lisp.

Emacs also acts as a shell, a batch manager, replacing startx / xinitrc, and a window manager . To control the execution of services, runit from the busybox package is used. There is an intention to rewrite runit and sinit to Lisp and run Emacs as PID 1.

SystemE-based environments can use packages from Kiss Linux , a minimalist distribution whose developers, in accordance with the KISS principle, are trying to build an extremely simple system, free from complications. The full-time package manager in KISS is written in shell and includes about 500 lines of code. All packages are compiled from source. Dependency accounting and overlaying additional patches are supported. Package metadata is placed in text files and can be parsed by regular unix utilities. Musl is used as the system C library, and the set of utilities is based on busybox. The initialization scripts are really simple.

<https://github.com/a-schaefer/systemE>

Qubes 4.0.3 OS Update Using Virtualization to Isolate Applications:

An update of the Qubes 4.0.3 operating system released , realizing the idea of using a hypervisor to strictly isolate applications and OS components (each class of applications and system services work in separate virtual machines). To work, you will need a system with 4 GB of RAM and 64-bit Intel or AMD CPUs supporting VT-x technology with EPT / AMD-v c RVI and VT-d / AMD IOMMU, preferably Intel GPUs (NVIDIA and AMD GPUs are not well tested) Templates for creating virtual environments based on Fedora 30, Debian 10 and Whonix 15 are available.

<https://www.qubes-os.org/news/2020/01/23/qubes-4-0-3/>

Sway 1.4 user environment release using Wayland:

Composite manager Sway 1.4 (no release 1.3) was released. It was built using the Wayland protocol and is fully compatible with the i3 mosaic window manager and the i3bar panel . The project code is written in C and distributed under the MIT license. The project is intended for use on Linux and FreeBSD.

Compatibility with i3 is ensured at the basic level, which allows you to use Sway as a transparent replacement for i3, using Wayland instead of X11 of course. Sway allows you to place windows on the screen as i3 would. The windows are arranged to form a grid that optimally uses the screen real estate and allows you to quickly manipulate the windows only using the keyboard.

<https://github.com/swaywm/sway/releases/tag/1.4>

SQLite 3.31 DBMS Released with Support for Generated Columns:

SQLite 3.31.0 is a lightweight DBMS designed in the form of a plug-in library. SQLite code is distributed as a public domain, i.e. can be used without restrictions and free of charge for any

purpose. Financial support for SQLite developers is provided by a specially created consortium, which includes companies such as Adobe, Oracle, Mozilla, Bentley and Bloomberg.

<https://www.sqlite.org/changes.html>

Android-x86 9.0-rc2 build available:

The Android-x86 project , in which an independent community is developing the port of the Android platform for the x86 architecture, published the second test release based on the Android 9 platform . It includes fixes and additions that improve the work of Android on x86 architecture. There are universal live-builds of Android-x86 9 for x86 32-bit (725 Mb) and x86_64 (920 Mb) architectures suitable for use on standard laptops and tablet PCs available for download . There are also rpm packages for installing the Android environment in Linux distributions.

<https://www.android-x86.org/releases/releasenote-9-0-rc2.html>

Good news for ChromeOS users:

Google has announced an extension of up to 8 years, for Chromebook device maintenance, including automatic updates. Initially, automatic updates for the Chromebook was three years, but then the support time was extended to six years, and now to eight. The reason for the support extension is the desire to extend the life cycle of equipment in schools that use Chromebooks in computer labs and reduce electronic waste.

<https://www.blog.google/outreach-initiatives/education/2020-chromebooks/>

GhostBSD Release 20.01:

GhostBSD 20.01 , based on the TrueOS platform and offering the MATE user environment, is available for download. By default, GhostBSD uses the OpenRC initialization system and the ZFS file system. It supports both live mode operation and installation on the hard drive (using the own ginstall installer,

written in Python).

The new version synchronized with the TrueOS 12.1-STABLE branch, the MATE desktop was updated to version 1.22. Fixes in the installer related to working with GPT and UEFI partitions were introduced. There is an added warning about using an incorrect password in Software Station. The ‘doas’ code was removed from NetworkMgr and replaced with the ‘sudoers.d / networkmgr’ file, to perform operations with elevated privileges.

http://www.ghostbsd.org/20.01_release_announcement

GameMode 1.5, a Linux game performance optimizer, is available:

Feral Interactive has released GameMode 1.5, implemented as a background optimization process that on-the-fly modifies various Linux system settings to maximize gaming application performance. The project code is written in C and comes under the BSD license.

In release 1.5, the ability to dynamically change the CPU governor for Intel processors with an integrated GPU has been added. Switching to “powersave” mode allows you to reduce the power consumption of the CPU and free up more GPU resources. The example touted is on the i7-1065G7 CPU, there the optimization allows you to increase the performance of “Shadow of the Tomb Raider” by as much as 25-30%.

<https://github.com/FeralInteractive/gamemode/releases/tag/1.5>

OASIS Technical Committee Approves OpenDocument 1.3 Specification:

The OASIS Consortium Technical Committee has approved the final version of the ODF 1.3 (OpenDocument) specification. After approval by the technical committee, the ODF 1.3 specification received the status of “Committee Specification”, which implies completion of the work, the specification and the readiness of the document is out for use by third-party

developers and companies.

The key difference between OpenDocument 1.3 and the previous version of the specification was the inclusion of new features to protect documents, such as document verification with digital signature and encryption of content using OpenPGP keys. The new version also includes clarification of the wording and some of the already available features were expanded.

<https://blog.documentfoundation.org/blog/2020/01/21/odf-1-3-approved-as-oasis-committee-specification/>

Canonical offers Anbox Cloud, a cloud-based platform for launching Android applications:

Canonical introduced Anbox Cloud, which allows you to run applications and play games created for the Android platform on any other system. Applications are launched on servers using the Anbox open environment, with streaming output to the client systems and transmitting events from input devices with minimal delays.

Ubuntu 18.04 LTS, LXDE, Juju and MAAS are used to organize the execution and orchestration of launching applications in containers. The components of the platform are developing as open projects, but in general, Anbox Cloud is a commercial product and is only available after filling out an application. The solution is optimized for servers based on ARM and Intel (x86) chips, and also supports graphics accelerator cards, such as the Intel Visual Cloud Accelerator Card.

Companies can use Anbox Cloud to transfer applications to public or private cloud platforms, making it possible to run them on any system without being tied to mobile devices. Game developers can use Anbox Cloud to expand their gaming audience. Mentioned were: organizing game streaming services (Game streaming), providing access to applications through the cloud, creating virtual devices, organizing work with corporate mobile applications, testing mobile applications (emulation of various types of devices is supported).

<https://ubuntu.com/blog/implementing-an-android-based-cloud-game-streaming-service-with-anbox-cloud>

LibreOffice 6.4 released:

LibreOffice 6.4 is now live on all supported platforms, bringing a long list of new features, but also performance optimizations and further compatibility improvements for Microsoft Office document formats. There are a lot of new features, the standout one being that you can now embed QR codes. Version 6.4 is now available on the LibreOffice website for download, if you can't wait.

<https://news.softpedia.com/news/libreoffice-6-4-released-with-new-features-performance-improvements-529024.shtml>

A new messaging client, Dino has been released:

This is the first release of the Dino app, that supports messaging using the Jabber / XMPP protocol and it is available for download. The program is compatible with various XMPP clients and servers. It is focused on ensuring confidentiality of negotiations AND supports end-to-end encryption using the OMMPO XMPP extension based on the Signal protocol or encryption using OpenPGP. The project code is written in Vala using the GTK toolkit and is distributed under the GPLv3 + license.

<https://dino.im/blog/2020/01/dino-0.1-release/>

Kali Linux 2020 released:

With this release, there are a few big changes. Throughout the history of Kali (and its predecessors BackTrack, WHAX, and Whoppix), the default credentials have been root/toor. The developers are no longer using the superuser account, root, as default in Kali 2020.1. The default user account is now a standard, unprivileged, user.

As with previous releases, it is not recommended to install Kali or use it as your main OS. All the desktop environments are

now contained in one ISO image. As Python 2 is end-of-life, the tools that depend on Python 2 have been removed.

<https://www.kali.org/releases/kali-linux-2020-1-release/>

FreeNAS 11.3 Network Storage Distribution Release

IXsystems released FreeNAS 11.3, a distribution kit for rapidly deploying network attached storage (NAS, Network-Attached Storage). The distribution is based on the FreeBSD code base. It features integrated ZFS support and the ability to manage the NAS via a web interface, built using the Python Django framework. FTP, NFS, Samba, AFP, rsync and iSCSI are supported for accessing the storage. Software RAID (0,1,5) can be used to increase storage reliability and LDAP / Active Directory support is implemented for client authorization.

<https://www.ixsystems.com/blog/truenas-and-freenas-11-3-release-pr/>

Heptapod public hosting announced for open source projects using Mercurial

The developers of the Heptapod project, which develops a fork of GitLab Community Edition, has moved to the Mercurial source code management system. They have announced the introduction of public hosting for Open Source projects (foss.heptapod.net) that use Mercurial. Heptapod code, like GitLab, is distributed under the free MIT license and can be used to deploy similar code hosting services on its servers.

The launched service allows free placement of any free and open projects using licenses approved by OSI. There is one condition - displaying the Heptapod sponsor logos (Clever Cloud and Octobus) on the official project web page. Due to the termination of support for Mercurial by Bitbucket hosting, priority will be given to applications from projects hosted on Bitbucket.

<https://heptapod.net/>

XCP-ng, a free version of Citrix XenServer, is now part of the Xen project

XCP-ng, is a free replacement of the proprietary cloud infrastructure XenServer management platform (Citrix Hypervisor). By joining the project the Xen, the development is now hosted in the Linux Foundation. Going under the mantle of the Xen Project will allow XCP-ng as standard distribution kit for deploying a virtual machine infrastructure based on the Xen and XAPI hypervisor.

XCP-ng virtualization platform is a Xen Project incubation project hosted in the Linux Foundation. XCP-ng aims to be the bridge between the users community and the developers and is the default entry point for any user. That will ensure the project will continue to develop independently. (It will not turn into a limited commercial product, as with XenServer).

<https://xcp-ng.org/>

Godot 3.2 Open-Source Game Engine Released With Better Documentation, New Features:

Godot 3.2 is providing better documentation, Mono / C# integration, Android (WIP) and WebAssembly, Oculus Quest support, overhauling of Godot's visual shaders, various graphics/rendering improvements, glTF 2.0 3D asset support, WebSocket and WebRTC support, new editor features, and a variety of other enhancements. Godot's editor is the main interface for the engine, and thus it got a ton of attention from all contributors. Here are some of the changes.

It is now possible to disable editor features. This allows to hide features that you don't intend to use to simplify the interface. This can be useful for tutors or companies who might want to restrict the access to some areas of the editor to let their students/teams focus on a specific subset.

Godot 3.2 offers up some modest improvements to its OpenGL-based code. Godot 3.2 brings improvements to its renderer

pipeline, various OpenGL ES 3.x features back-ported to GLES2 like MSAA, and enhancing the default renderer settings. It is available for download now.

[*https://godotengine.org/article/here-comes-godot-3-2*](https://godotengine.org/article/here-comes-godot-3-2)



DistroWatch.com

Put the fun back into computing. *Use Linux, BSD.*



The Ubuntu Podcast covers all the latest news and issues facing Ubuntu Linux users and Free Software fans in general. The show appeals to the newest user and the oldest coder. Our discussions cover the development of Ubuntu but aren't overly technical. We are lucky enough to have some great guests on the show, telling us first hand about the latest exciting developments they are working on, in a way that we can all understand! We also talk about the Ubuntu community and what it gets up to.

The show is presented by members of the UK's Ubuntu Linux community. Because it is covered by the Ubuntu Code of Conduct it is suitable for all. The show is broadcast live every fortnight on a Tuesday evening (British time) and is available for download the following day. podcast.ubuntu-uk.org

[RETURN TO CONTENTS](#)

Command & Conquer

Written by Erik

With Lucas “retiring”, I thought we could fill his space with something that is more of the same, yet different.

I want to vault us into some ‘capture the flag’ type-of-stuff, but we need to lay some groundwork first. I will use the words directories and folders interchangeably. The reason for this is that there are GUI jockeys out there who will be reading this, who are not interested in terminal (shell) terminology. I want them to feel welcome reading this piece.

Let’s talk about privilege escalation in Linux. Because Linux has an all powerful user named ‘root’, that is our aim, well... a root shell... Sometimes getting root is as easy as exploiting a program or misconfigured service. Sometimes you have to string things together to step over the wall, so to speak. Sometimes you need to push the enter key 790 times.... But I am getting ahead of myself. We need to talk about permissions first. (What is privilege escalation, but breaking and defying your set permissions?) <play “i = "" want = "" to = "" break = "" free” = "" by = "" freddy = "" mercury = "" for = "" dramatic = "" effect = "" >:)

When you look at Linux permissions, there are users and groups on one side, and files and folders (directories) on the other. (Yes, ‘others’ too, but I want you to paint a mental picture, OK?)

All users’ passwords are hashed and stored in “/etc/shadow”, but the accounts are still in “/etc/passwd”. You will notice root has a UID of 0, and, depending on your Linux system, users will start at UID 500 or UID 1000. UID 0 has access to everything, thus making it the target for privilege escalation. Groups can have many users, and users can be in many groups. However, all files and folders have one owner and group. Like the other

permissions, groups are stored in `“/etc/groups”`. Permissions granted or denied are by read (r), write (w), and execute (x). Do a `‘ls -la’` in the folder where you are right now. Look at the leftmost column. You will likely see something like `“-rw-r--r--”` a number, then the owner (user and group). Only the owner or root can change permissions. The very first one is sometimes `“d”` for directory, so ignore that. Thereafter, the first triplet is for the owner, the next triplet is for the group, then the last triplet is for `“other”`. `“Other”` is sometimes called `“world”`, but do not confuse the (w) for `‘write’` with `‘world’`.

On a file (r) (w) or (x) is just what they say, but they are slightly different on directories / folders. Here, (r) is list contents, (w) allows sub-directories to be created, and (x) is the gatekeeper, allowing you access to the other two permissions. If (x) is not set, you cannot get to (r) or (w).

But it does not end there. There are SUID and SGID bits. When this bit is set, SUI, files execute with the permissions of the file owner. The same with SGID. SGID also allows sub-directories to inherit group permissions from the parent. You will see an (s) in the (x) position.

Which brings us to `“su”` and `“sudo”` and the like. When you change users, you get what is known as an `‘effective ID’`, which is not your `‘real ID’`. (Thank goodness there are no step-permissions! Hahaha!). The effective ID is used to verify your access more often than the real ID. Now if that was not confusing enough, there is also a `“saved ID”`. This keeps track of who you really are, when you switch permissions temporarily.

When you use the `“who”` command, you get your effective ID, and when you use the `“id”` command, you should see your real ID as well as your effective ID. Look for (euid) and (egid). Most of the time, these will be the same as your UID. Run the `“id”` command as yourself, then run it as `“sudo id”`, and `su` before running `id` and see how it changes.

Processes also have permissions. To see the permissions of your terminal (shell), type: `‘cat /proc/$$/status | grep id’` and in the

centre you will see, Uid and Gid. Note that the shell is the current process. If your system is Ubuntu, you should see 1000 repeated a couple of times, and if you used su, you should see 1000 followed by zeroes. (I am assuming that you are the only user on your system). You also need to know that Linux is divided into user space and kernel space. Even your CPU is divided into user mode and protected mode. Are you seeing a pattern here? On the CPU, you usually have “rings” – ring 0 through ring 3 if I remember correctly. (Hello Greg!, I sort of remembered!). However, in modern CPUs, they go into the negative too. Linux uses ring 0 for kernel space and ring 3 for user space. This is why some hackers work so hard on “undocumented features” because, if you can talk to ring 0, you can bypass all the operating system permissions. This is beyond what we are going to cover, but at least you know about it.

Did you learn something today? Did I make a boo-boo? Are you as excited as I am? Let us know. misc@fullcirclemagazine.org

RETURN TO CONTENTS

How To - Python in the REAL World - Part 101

Written by Greg D. Walters

The other morning, while I was getting my shower, my mind went to a fairly dark place as it often does. If most people fall in the shower, or outside, or going down the stairs, they will end up with a bruise and be achy for a day or two. If I fall however, there is a very high probability that I will be paralysed or worse. And, as usual, when I get into that place, I wonder how I'll be able to continue my writing, and programming, and cooking, if that happened. I'm sure that I'm not the only one who thinks about that kind of thing.

Luckily, today we have Siri, Amazon Alexa, Google Assistant, and more. Almost every smartphone has some kind of speech recognition. There are many pre-made packages out there for Linux and other operating systems. But I wanted to see what could be done via Python.

First, I want to hit the pause button and share a little history of speech recognition with you. Back when I was a child, when rainbows were in black and white, and we had to watch TV by candlelight, because there wasn't any electricity (not really, but it confuses children to no end), computers were just getting started. In 1952, Bell labs created the Audrey system, which was able to understand a single speaker speaking numbers. Fast forward 10 years and IBM created a system called "Shoebox" which could understand and respond to a whopping 16 words. (see <https://sonix.ai/history-of-speech-recognition>)

Enough of ancient history. Push the play button!


After a little web browsing, I found a library for Python called, surprisingly enough, SpeechRecognition. It can be installed via pip...



All the source code can be found at https://github.com/Uberi/speech_recognition#readme .

I went ahead and installed via pip, and then went and downloaded the source from the github repository.

I've "borrowed" the following snippet from the repository site...



Now, there are somethings that need to be said here. Most of these online engines require you to register as a user and obtain keys to be able to access them, and/or may incur costs. The only offline services that are currently supported are CMU Sphinx (I'll talk about that in a little bit) and Snowboy. If you wish to see exactly what is required for each engine, download the source code from the GitHub repository and look at the **init.py** file located in the `speech_recognition` folder of the distribution.

Once I saw the line "To quickly try it out, run `python -m speech_recognition` after installing.", I couldn't resist. But I did, at least long enough to see what other requirements there might be, and I'm glad I did. A little bit further down says that

if you want to use a microphone, which of course I do, that you need to install PyAudio. Ok. That makes sense. So I read a little further on. I saw this...

“On Debian-derived Linux distributions (like Ubuntu and Mint), install PyAudio using APT: execute `sudo apt-get install python-pyaudio python3-pyaudio` in a terminal.”


I immediately copied the `apt-get` line from the bullet point and ran it in a terminal. I noticed that I could use `pip` to install the actual library. HOWEVER, because I was being stupid, I didn't notice the caveat below that.

“If the version in the repositories is too old, install the latest release using Pip: execute `sudo apt-get install portaudio19-dev python-all-dev python3-all-dev && sudo pip install pyaudio` (replace `pip` with `pip3` if using Python 3).”


When I tried to use `pip`, I got a tonne of error messages. That's because I had not installed `portaudio19-dev` with the rest of it.

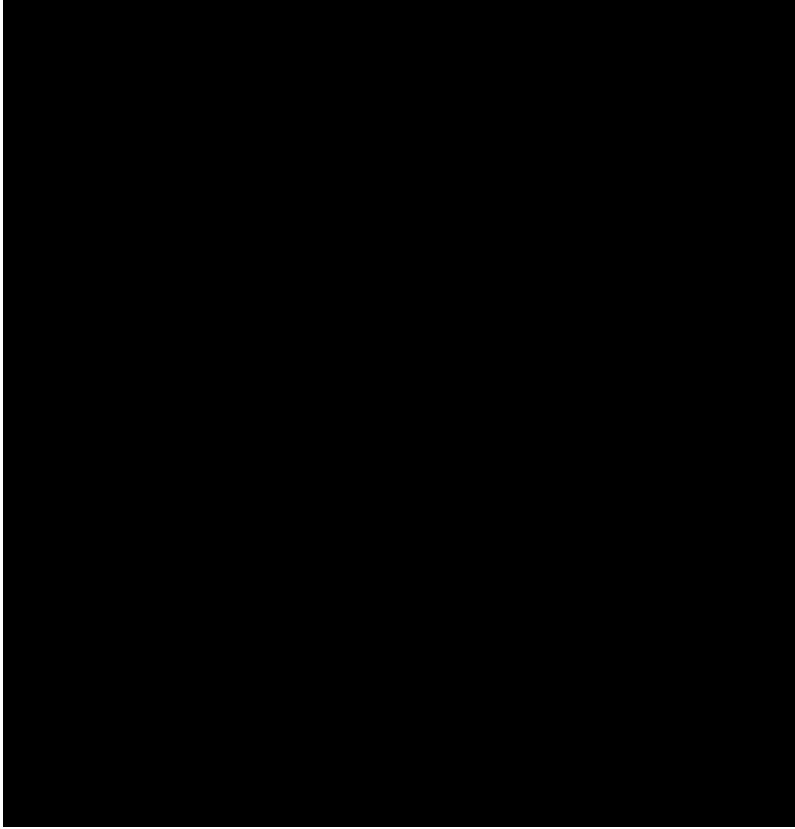
I did another `apt` install, and ran the `pip` install command. It worked!

So to get it all on your system, here's what you want to do...




Now, you can give the program a try...



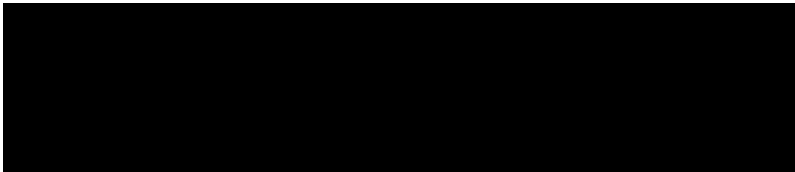


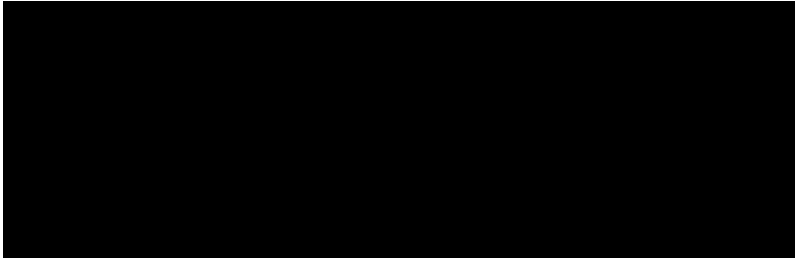
Ok. Color me impressed. The warning messages didn't worry me, they actually piqued my interest about all the possibilities. I wasn't really happy with having to hit <Ctrl> <C> to get the program to quit though. Now, I wanted to know more.

Digging into the distribution folder, I found **main.py**, which I quickly modified to put my own unique spin on it...

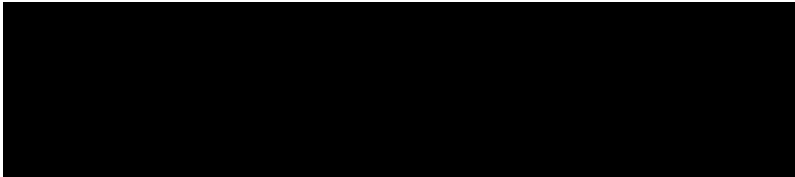


First, we import the library and create instances of the Recognizer and the microphone objects. Next, we use a loop to continually check and adjust the microphone level...

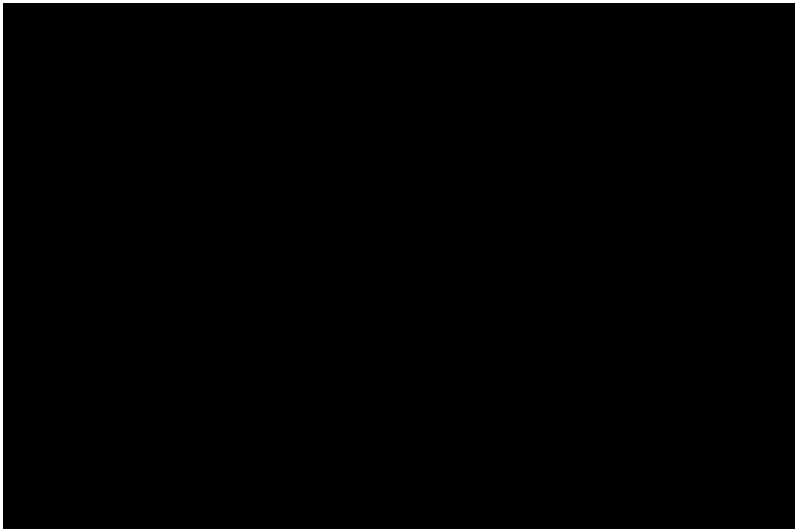




Now, we create another loop to get something that looks like speech from the microphone...



This continues to listen until a) it hears speech, and b) the speech stops. Then it tries to process the speech audio...



Notice that in this step, the Recognizer uses the Google Speech Recognition system. The line “if str is bytes:” checks to see if this is running under Python 2.x to properly print any unicode characters. Now we have printed out what the Recognizer THINKS was said. Next, we can check that with either the phrase “please quit” or “please stop” to programmatically end the program. I tried using a single word, but that never

triggered. I'm guessing that the system just figured that the input was just noise. A side note here. What if you don't want to use English as the language that you speak to the program with. What about Spanish or Norwegian or some other language. It's covered! Change the line:

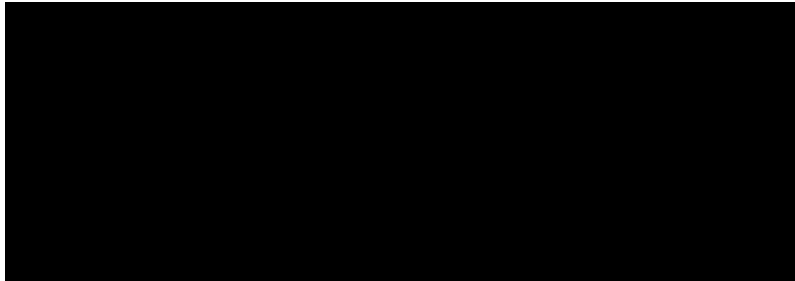


To

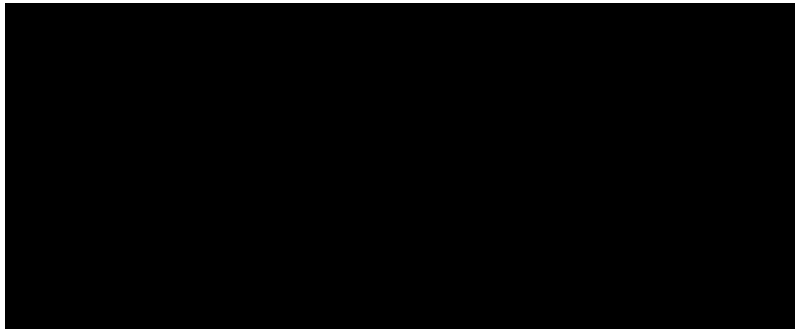


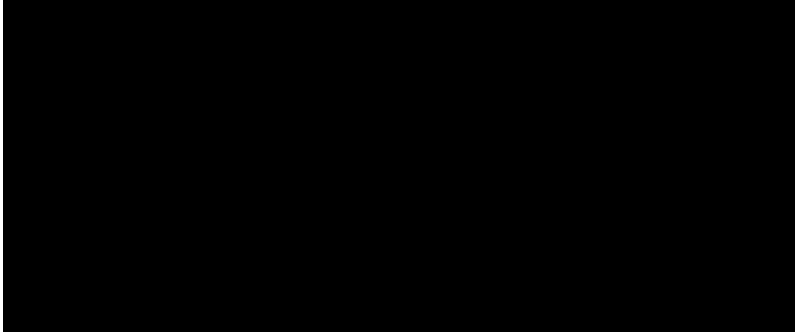
For British english, "no-NO" for Norwegian, or "es-AR" for Spanish (Argentina). You can check this link (<https://stackoverflow.com/questions/14257598/what-are-language-codes-in-chromes-implementation-of-the-html5-speech-recogniti/14302134#14302134>) to see many language settings. Moving on...

Finally, we check for exceptions...



It looks VERY easy, doesn't it. It is. Here's a sample of the program run (I removed the warning messages from the output)...



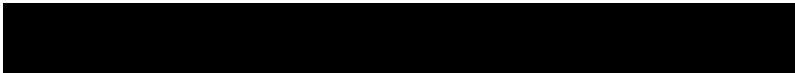


So the bottom line is that we get a string (value) back from the Recognizer. What we do with that information, right now, is open ended. A friend suggested that it might be good to use the Text to Speech API with `espeak-ng` that I talked about back in Full Circle Magazine #150. We might revisit in a future article.

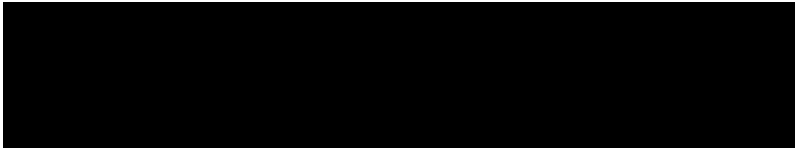
At this point, before I forget it, I promised a while ago to talk about PocketSphinx. There are a lot of people who find that it is not very reliable. I tried to get it to install and I have to admit, there were issues.

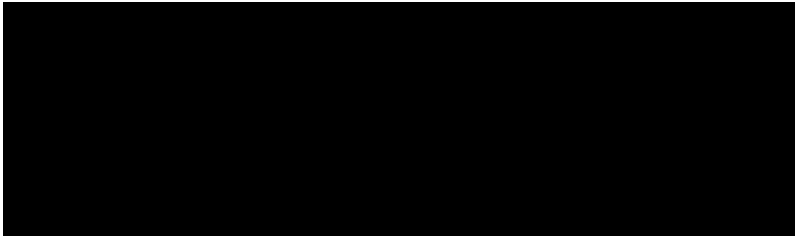
First, I suggest that if you want to try PocketSphinx, you go to <https://pypi.org/project/pocketsphinx/> and follow the instructions there. The GitHub repository is at <https://github.com/bambocher/pocketsphinx-python> .

There is an example program that is provided with the source distribution that, at least for me, would not run. I kept getting an error starting with the line `decoder = Decoder(config)`. I did a search and found a number of people having the same issue, but not much in the way of an answer. After digging for much longer than I should have, I found a reference to the `MODELDIR` config settings. After looking into my Python library folders, I found the site package for PocketSphinx. I realized that the `MODELDIR` and `DATADIR` statements were not being set properly in the example. They were:



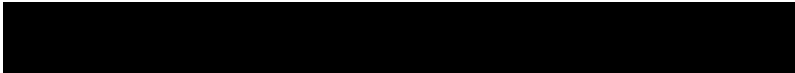
but for me, they needed to be...



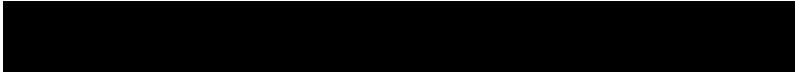


The package needs absolute path statements to where pip installed PocketSphinx. This can be a major problem if you are using something like pyenv and have multiple instances of Python or if you wish to distribute an app you wrote using PocketSphinx.

The next problem showed up on the line:



This was incorrect based on the installation. The “files are located in a folder directly off of the model folder. It should have been:



After these changes were made and saved, the example program worked.

You might wonder at this point, “Ok, so how do we actually do something with the data we’ve received?” That is such an open-ended question, that it’s really out of the scope of this article. HOWEVER, I can point you in an interesting direction.

If you remember near the top of the article, one of the engines that is supported by the SpeechRecognition library is Wit.ai . This is an interesting site. Basically, you provide speech or text to their API and it tries to match that input to something you have told the system that you expect the end user to enter. For example, let’s say that you want your end user to say things that would be along the lines of home automation like turning on or off a light, asking what the temperature is outside, changing the thermostat, and so on.

Check out <https://wit.ai/>. It takes a little bit of navigation, and you have to read a bunch of the site to understand, but I think you'll get the gist of things pretty quickly. We'll explore it some more next time.

One other thing. While I was digging around on the web to get info to do this article, I found that Google Chrome now can support voice commands. I haven't tried it yet, but it looks very interesting. From the website, it says "Use the magic of speech recognition to write emails and documents in Google Chrome. Dictation accurately transcribes your speech to text in real time. You can add paragraphs, punctuation marks, and even smileys using voice commands." Check out the site at <https://dictation.io/>

The code is, as always, on Pastebin at: <https://pastebin.com/pTJ6RcKL>

Until next time, keep coding!

Greg Walters is a retired programmer living in Central Texas, USA. He has been a programmer since 1972 and, in his spare time, he is an author, amateur photographer, luthier, fair musician and a pretty darn good cook. He still is the owner of RainyDaySolutions, a consulting company, and he spends most of his time writing articles for FCM and tutorials. His website is www.thedesignedgeek.xyz.

RETURN TO CONTENTS

How-To - Test Linux in VirtualBox

Written by Chris Christensen

I have been using the Linux operating system for nearly twenty years. Those twenty years have not always been smooth, and that is probably why Linux hasn't gained the traction that Windows or MacOS have in the market. In the past, Linux took some work to get up and running. There are plenty of horror stories of getting wifi working on laptops. I have told those stories!

Times have changed. Red Hat Linux was recently purchased by IBM for \$34 billion. That is a lot of zeroes for a company that was built on an operating system that is freely available to anyone that wants it. Red Hat Linux will not be the focus of this article as it is more of an "Enterprise" operating system. Ubuntu Linux is a much better personal operating system, and I run it on my personal laptop. Ubuntu has become so popular and user friendly that Dell and Lenovo both offer systems with Ubuntu Linux pre-installed. For our purposes, the Ubuntu operating system is easy to download and install.

Back in my day (slight grin at my aging rhetoric), you would have to install Linux on to a physical computer to give it a try. You could dual boot with Windows but that can be a painful process for new users. You can try Ubuntu Linux out as a "Live" system without installing but the concept may be lost on people that are not as familiar with technology. The most painless approach is to install Linux on a virtual machine using VirtualBox. Like Ubuntu Linux, VirtualBox is easy to download and install.

Let's get started!

VirtualBox Installation

Go to <https://www.virtualbox.com> and click on the large green button on the homepage.

On the next page you will see a heading like this:

VirtualBox 6.0.10 platform packages

- ➞ Windows hosts
- ➞ OS X hosts
- Linux distributions
- ➞ Solaris hosts

On the same page you will see a heading like this:

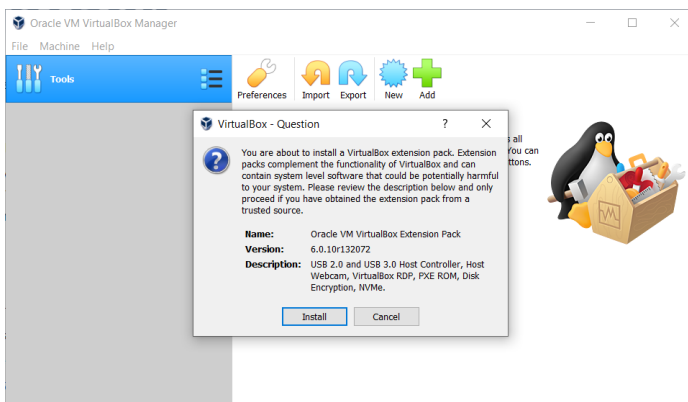
VirtualBox 6.0.10 Oracle VM VirtualBox Extension Pack

- ➞ All supported platforms

Download the extension pack as well.

Once you have VirtualBox downloaded, launch the downloaded file to begin the installation. You can click next all the way through as the defaults are fine.

When VirtualBox is completely installed, open the VirtualBox Extension Pack file that you downloaded and install it as well.

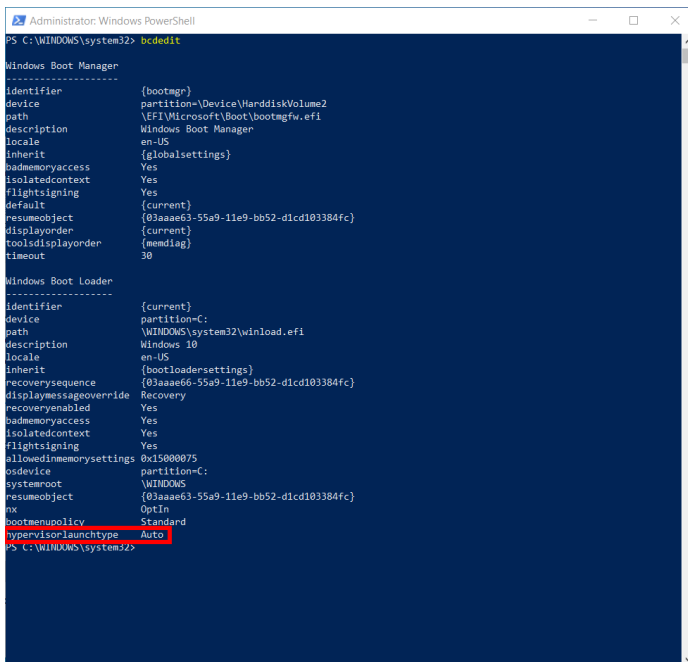


Hypervisorlaunchtype Issue on Windows 10

At this point, you could attempt to create your first virtual machine but you may run into the same issue that I ran into on Windows 10. Windows 10 has a setting called “Hypervisorlaunchtype”. This obscure setting will block you from being able to start your virtual machine on VirtualBox. I found a blog post from Zahid Anwar that outlines how to disable this feature.

Open Powershell as an administrator. If you are not familiar with Powershell, it is a windows terminal program that you can find by clicking on your start menu button and typing “powershell.” When it appears in the menu, right-click on the icon and choose “Run as Administrator.”

With Powershell open, type `bscdedit` at the prompt.



```
Administrator: Windows PowerShell
PS C:\WINDOWS\system32> bcdedit

Windows Boot Manager
-----
identifier              {bootmgr}
device                  partition=\\Device\\HarddiskVolume2
path                    \\EFI\\Microsoft\\Boot\\bootmgr.efi
description              Windows Boot Manager
locale                  en-US
inherit                  {globalsettings}
badmemoryaccess          Yes
isolatedcontext          Yes
flightsigning            Yes
default                  {current}
resumeobject             {03aaae63-55a9-11e9-bb52-d1cd103384fc}
displayorder             {current}
toolsdisplayorder        (nvidia)
timeout                  30

Windows Boot Loader
-----
identifier              {current}
device                  partition=C:
path                    \\WINDOWS\\system32\\winload.efi
description              Windows 10
locale                  en-US
inherit                  {bootloadersettings}
recoverysequence         {03aaae63-55a9-11e9-bb52-d1cd103384fc}
displaymessageoverride   Recovery
recoveryenabled          Yes
badmemoryaccess          Yes
isolatedcontext          Yes
flightsigning            Yes
allowedinmemorysettings 0x15000075
osdevice                 partition=C:
systemroot               WINDOWS
resumeobject             {03aaae63-55a9-11e9-bb52-d1cd103384fc}
nx                       OptIn
bootmenupolicy            Standard
hypervisorlaunchtype     Auto
PS C:\\WINDOWS\\system32>
```

The part I have highlighted in red is what you are looking for. If this is set to “Auto” you will need to run `bscdedit /set hypervisorlaunchtype off`.

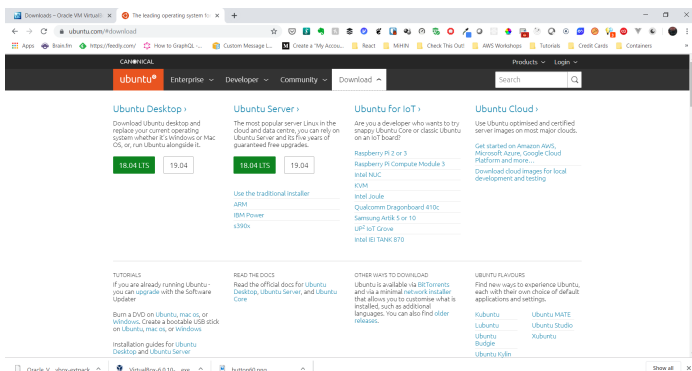
image

After you have turned off `hypervisorlaunchtype` you will need to reboot your computer.

Download Ubuntu Linux

Before we create our first virtual machine we need to download Ubuntu Linux from the Ubuntu website. Go to <https://www.ubuntu.com>

Click on Download in the navigation.



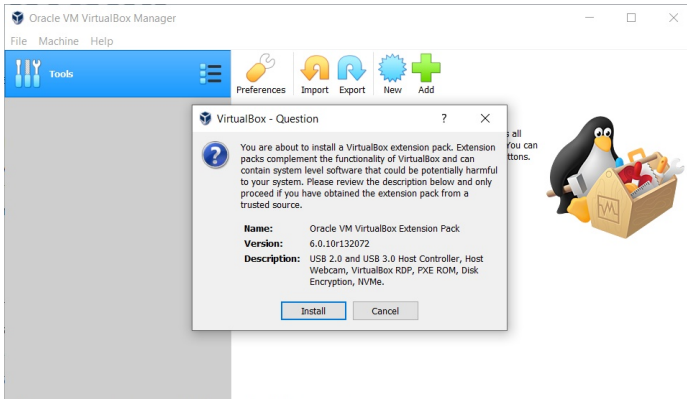
Under the “Ubuntu Desktop” heading choose “18.04 LTS” or a later LTS version. LTS stands for “Long Term Support” and is a version of Ubuntu that is supported for five years after it is released.

Clicking on its button will download the latest ISO image. This might take a few minutes depending on your Internet connection.

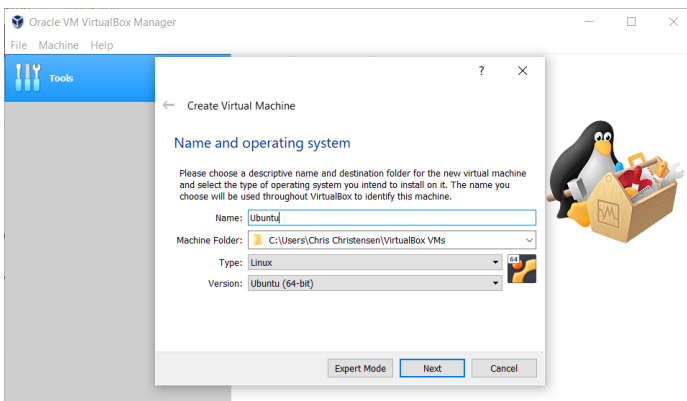
Create Your First Virtual Machine

Open VirtualBox

In the right-side panel above the “Welcome to VirtualBox!” heading are five choices. Click on the blue “New” button.



Enter a name for your virtual machine. This can be anything you want but with any naming convention make it meaningful. You never know. You might fall in love with creating virtual machines and spin up a dozen of them. Meaningful names help you tell them apart.



If you type in “Ubuntu”, VirtualBox will automatically change the type to “Linux” and the Version to “Ubuntu (64-bit).” If you use another name, make sure to click on “Type” and change it to “Linux” and “Version” and change it to “Ubuntu (64-bit).”

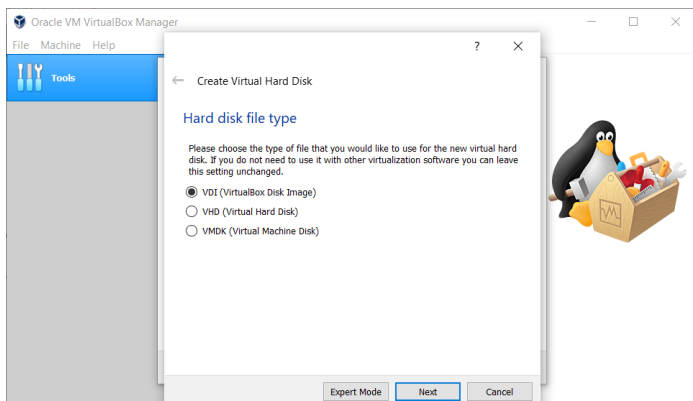
The next step is to choose how much memory you want your virtual machine to use. The recommended amount of 1024 MB is fine and can be changed later if you want to add more.

Next up is the disk for the virtual machine. Just like a physical

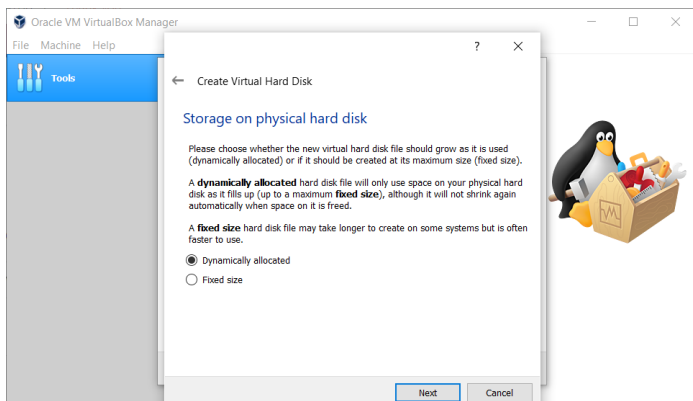
computer, a virtual machine needs a disk for storage.

Click “Create” to start the process of creating the virtual hard disk.

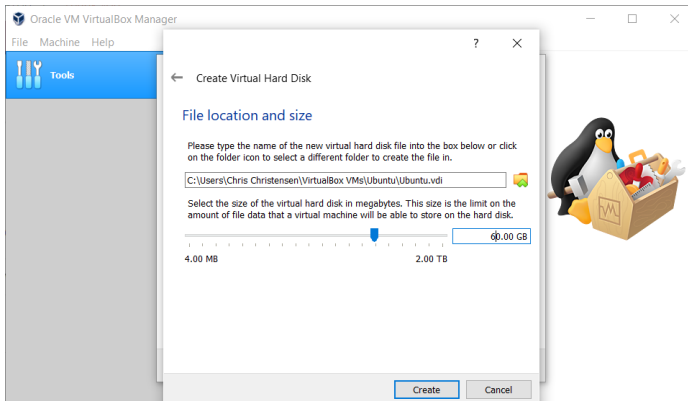
The next page will ask for the type of hard disk. The default “VDI” is fine.



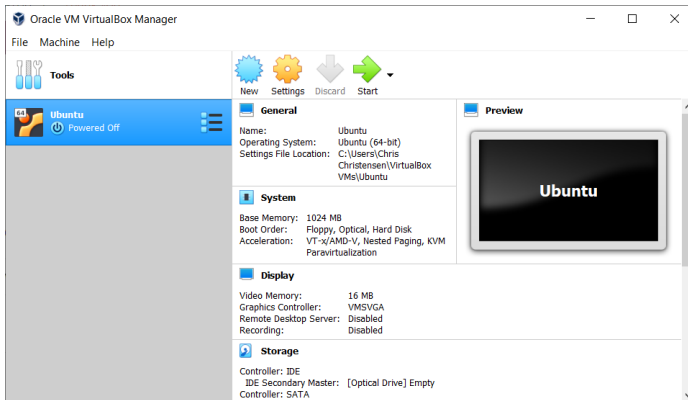
The next screen will ask if you want to use a “Dynamically allocated” or “Fixed size” disk. Dynamically allocated disks grow as you need them to grow. For instance, if you set the disk up with 80 GB of disk space but choose to dynamically allocate that disk space, VirtualBox will only use as much space as it needs until it hits 80 GB. On the other hand, fixed size will reserve the entire 80 GB from the beginning. The default dynamically allocated is fine and my preferred choice.



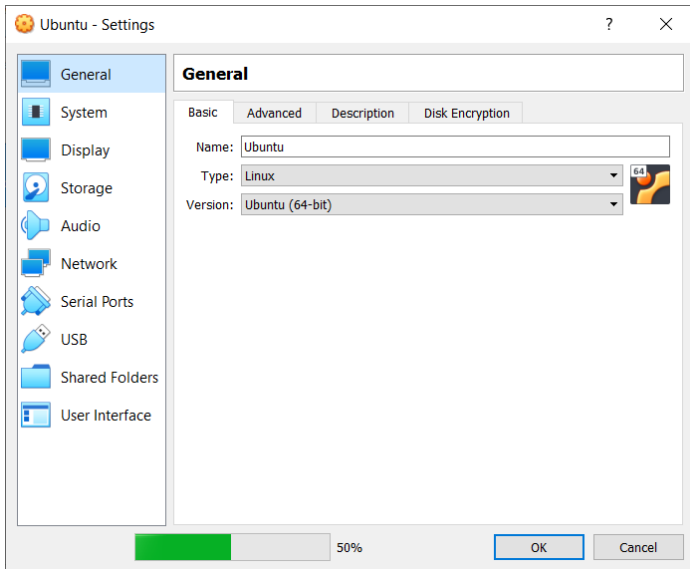
Next up is the size of the disk. You don't need a lot of disk space to take Ubuntu Linux for a test drive. I allocated 60 GB of disk space because I plan on using my virtual machine for something that requires this. Keep in mind that you can always delete the virtual machine and recover back any disk space that it takes up.



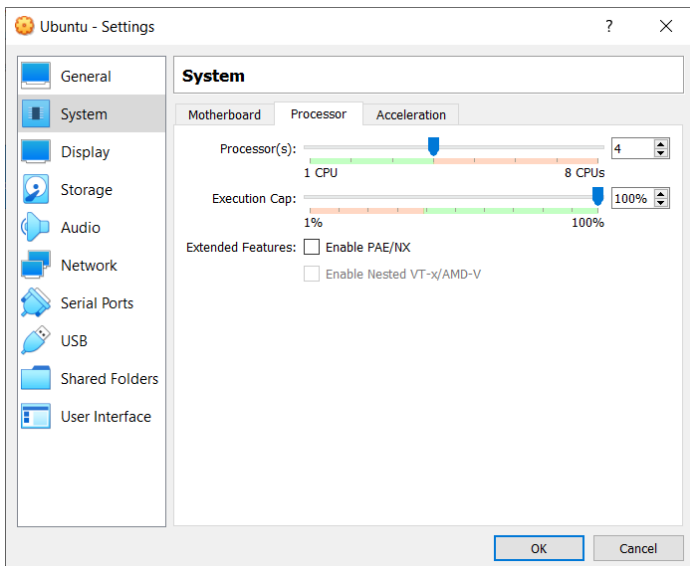
Once the disk is created you are returned to the VirtualBox main screen. Your new virtual machine should be highlighted. Click on the yellow “Settings” button in the right-hand panel.



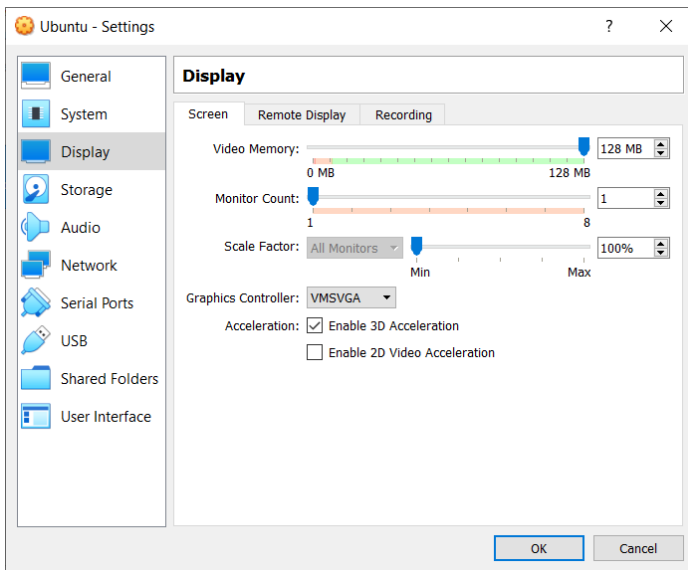
We are going to make a few changes to the virtual machine to optimize performance. We also need to tell the virtual machine to use the Ubuntu Linux ISO image that we downloaded earlier.



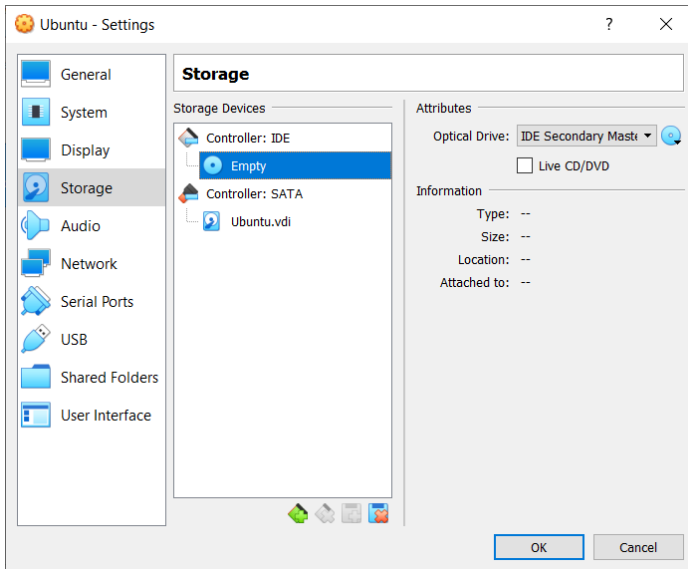
Click on “System” in the left-hand panel. Click on the “Processor” tab. Move the processor toggle to the end of the green. In my case, this was four CPUs. This is not necessary but will give you a bit more processing power with your virtual machine.




Next click on “Display” in the left-hand panel. Move the “Video Memory” toggle all the way to the end. This will give you 128 MB of video memory. You won’t be doing any hardcore gaming with this setting but it gives a bit better video performance.



Click on “Storage” in the left-hand panel and then click on the disk icon with “Empty” next to it.



In the attributes section in the right-hand panel click on the disk icon with the down arrow on it.

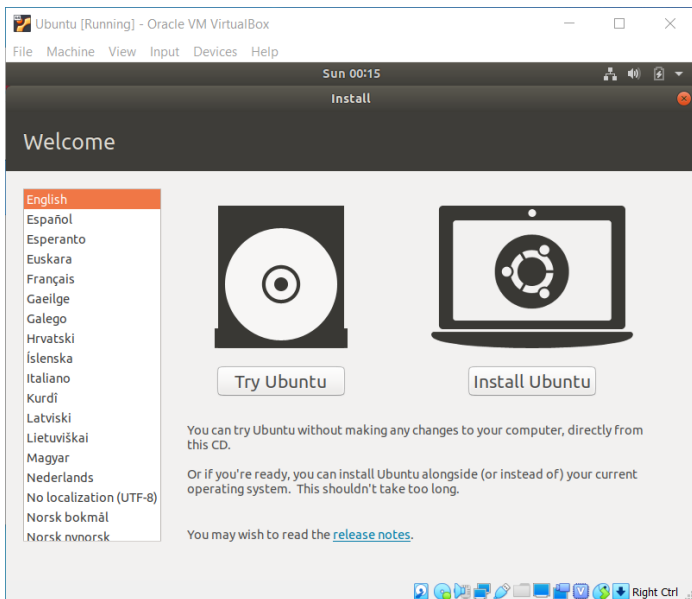


Note that my screenshot already shows “ubuntu-18.04.3-desktop-amd64.iso.” You will only see “Choose Virtual Optical Disk File...” when you click on this. Click on “Choose Virtual Optical Disk File...” and locate the Ubuntu Linux ISO image that we downloaded earlier.

Where it said “Empty” before, it now shows the image we downloaded.

We are done making changes to the configuration. You can click “OK.”

You should be back on the main screen for Virtualbox. Go ahead and double-click on your virtual machine to launch it. After an initial boot up, you should see the following screen.



Click on “Install Ubuntu.” On the next screen you can click continue unless you prefer a different language than English.

The default settings for “Updates and other software” are fine. Since this is a virtual machine there is no need to install third-party software for graphics and wifi.

We are not going to do any special partition and encryption and LVM (Logical Volume Manager) are probably overkill for a test drive so go ahead and click “Install Now.”


You will see a warning that if you continue the disk will be erased. It is safe to click “Continue.”

Next up is your time zone. You can click your location on the map or type in a location in the box that currently says “New York.”

The next screen is to set up your user account. You can put anything you’d like for these.


Finally the installation begins!

The installation will take a few minutes. When it is finished you will see a message to restart.



After clicking on “Restart Now”, you will see another message to remove the disk and click enter. There is no disk to remove so you can just hit enter.

After reboot, you should be presented with the Ubuntu login screen.



Initial Setup for Ubuntu Linux

You now have Ubuntu Linux installed and you could take it for a test drive. The problem is that it is still a tiny screen on the VirtualBox emulation screen. This isn't exactly optimal or appealing. We want to be able to experience Ubuntu in fullscreen, as if it is installed on our computer. To do this we need to install something called "Guest Additions."

Login to Ubuntu Linux with the password you setup previously.

You will see an initial set of screens on first login. You can click through these screens. The only screen of interest is the screen asking if it is ok to send information back to Canonical, the creator and maintainer of Ubuntu Linux. I choose to send the information back with the thought that they are using the information to make Ubuntu better. Privacy hawks will jump all over me for this. The choice is yours.

Most likely you will see a box that says that updated software is ready to be installed. Go ahead and install the update.

After the installation finishes you will be asked to reboot the

system. Go ahead and reboot.

We now need to install a package that is necessary to install Guest Additions. We need the terminal for this. Click on “Activities” in the upper left-hand corner and then type “terminal” in the search box. The only choice is the terminal program. Click on it to open the terminal.

At the prompt, type in

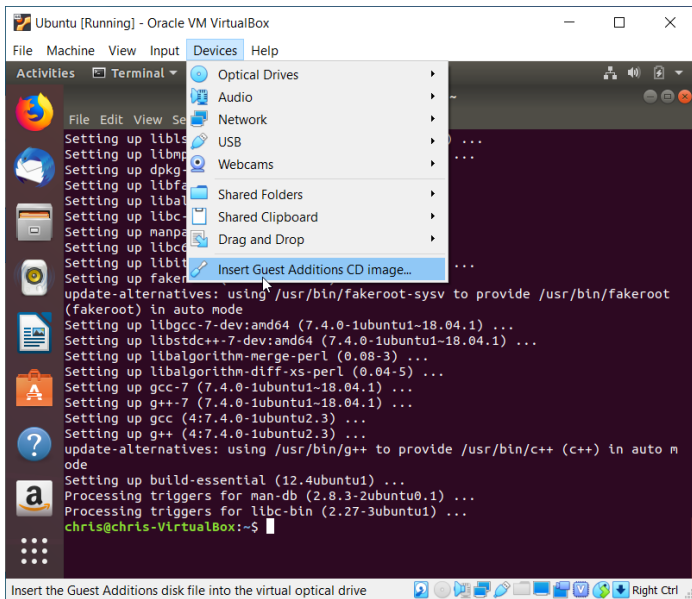


You will be prompted for your password. Type in the same password you use to login to the system.

You can just hit enter when prompted with the Y/n question.

```
Ubuntu [Running] - Oracle VM VirtualBox
File Machine View Input Devices Help
Activities Terminal Sat 20:56
chris@chris-VirtualBox: ~
chris@chris-VirtualBox:~$ sudo apt install build-essential
[sudo] password for chris:
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following additional packages will be installed:
  dpkg-dev fakeroot g++ g++-7 gcc gcc-7 libalgorithm-diff-perl
  libalgorithm-diff-xs-perl libalgorithm-merge-perl libasan4 libatomic1
  libc-dev-bin libc6-dev libcilkrts5 libfakeroot libgcc-7-dev libitm1
  liblsan0 libmpx2 libquadmath0 libstdc++-7-dev libtsan0 libubsan0
  linux-libc-dev make manpages-dev
Suggested packages:
  debian-keyring g++-multilib g++-7-multilib gcc-7-doc libstdc++6-7-dbg
  gcc-multilib autoconf automake libtool flex bison gcc-doc gcc-7-multilib
  gcc-7-locales libgcc1-dbg libgomp1-dbg libitm1-dbg libatomic1-dbg
  libasan4-dbg liblsan0-dbg libtsan0-dbg libubsan0-dbg libcilkrts5-dbg
  libmpx2-dbg libquadmath0-dbg glibc-doc libstdc++-7-doc make-doc
The following NEW packages will be installed:
  build-essential dpkg-dev fakeroot g++ g++-7 gcc gcc-7
  libalgorithm-diff-perl libalgorithm-diff-xs-perl libalgorithm-merge-perl
  libasan4 libatomic1 libc-dev-bin libc6-dev libcilkrts5 libfakeroot
  libgcc-7-dev libitm1 liblsan0 libmpx2 libquadmath0 libstdc++-7-dev libtsan0
  libubsan0 linux-libc-dev make manpages-dev
0 upgraded, 27 newly installed, 0 to remove and 0 not upgraded.
Need to get 26.9 MB of archives.
After this operation, 117 MB of additional disk space will be used.
Do you want to continue? [Y/n]
```

Now click on “Devices” in the VirtualBox menu and choose “Insert Guest Additions CD Image...”



A box will pop up asking if you want to run the software. Click “Run.”

You will be prompted for your password. Enter your password and click “Authenticate.”

Guest Additions will be installed. When it is finished you will be asked to press “Return” to close the window.

You will need to restart Ubuntu for the changes to take effect. Click on the small down-arrow in the upper right-hand corner and click on the power button which is the button at the bottom on the right-hand side.


Choose “Restart” from the options.

Once the system reboots and you log back in, expand the window by clicking on the box next to the “x” in the upper right-hand corner of the window.

Now click “View” in the VirtualBox menu and choose “Full-screen Mode.”

If you hover your mouse at the bottom-center of the screen, a menu will pop up that gives you the same options you used to have at the top of the screen before full-screen mode.

If you click on the square button next to the “x” you will be taken out of full-screen mode.



Conclusion

You are now ready to take Ubuntu Linux for a proper test drive. Sixty-one steps seems like a lot of work for a test drive but in the end you will get to experience Ubuntu Linux and the Linux operating system as it was intended.


If you are new to Linux I'd love to hear how your test drive went. If you are an experienced Linux user and have feedback or suggestions don't hesitate.

Chris Christensen is a Cloud Engineer at MiHIN. AWS, DevOps, Linux, Kubernetes, and all that good stuff. He is also an Accidental developer.

How-To - Darktable - Part 11

Written by Erik

If you have been following our Darktable tutorials, you will know we are working with the latest version, not the version that is bundled in the Ubuntu Software centre.



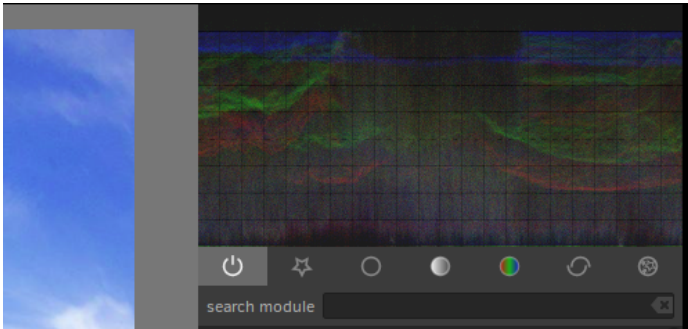
White balance, what is it? Well, it is the process of removing unrealistic color casts, so that objects which appear white in person are rendered white in your photo. Sounds simple, doesn't it? Whilst most of us play it by eye, there is a way to set your white balance without ogling the picture every five seconds. Truth!

What you need to know:

Though you can adjust the white balance on .jpeg-files, this is easier and more common to adjust it on RAW files. Because of the compression used in .jpeg-files, you adjust more than just the white balance afterwards.

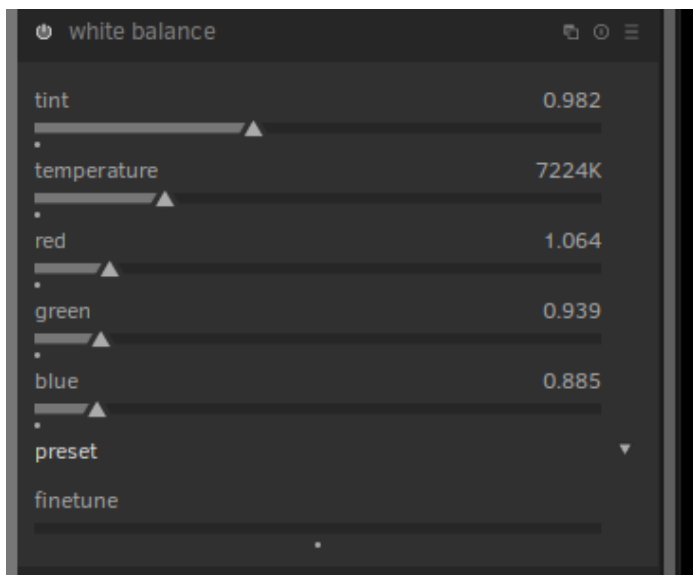
Before I ramble on about white balance, know that I live in sunny Africa, where clouds are few and far between. So my reference to white balance is usually the sun. You may live in an area where it is more cloudy, so the white balance in your photo may be an average of the light. The same goes for you guys in the city, with all its lights and reflective surfaces. (As well as night photography).

So the fastest way to set the white balance is to select a pure white, or neutral reference, in your photo. That's it, and I will show you how a bit later. The other way is to change your histogram in Darktable to waveform and try to align the colors. The two modules you will need are white balance and color correction.



If you look at my histogram here, you will see that blue is high and red is low. To get a more natural color, I need to raise the reds and lower the blues, without creating color casts. As this picture is a .jpeg-file, my chances of changing it much without making it unrealistic are slim. With RAW images, it is another ball-game altogether. This is why any photographer worth his salt shoots in RAW format. My father used to spend hours fiddling, setting his camera settings, before taking a photo. In effect, doing all this processing prior to getting a shot off. Let's call it pre-processing, as film was not really editable afterwards. This is the appeal of digital photography for me. I can go out, snap pictures of animals, before they run away, and fix the small things later. Also, Darktable allows me to save my presets, as, when taking photos up the skeleton coast, the environment and lighting seldom change. The sky is always blue, the sand always white, and the sea is always green. But it may be different for you, so white balance may be more of a "thing".

Open your picture in Darktable with the two modules that I mentioned. Go to your white balance module and you will see this:




You will see ‘preset’ with a triangle to the right. Expand that and change it to “spot”. This will allow you to select a spot in your picture that is as white as possible. That selection will direct Darktable’s algorithm to do an “auto white balance” for you. Click and drag a rectangle on anything white. If your photo contains no white, find a neutral spot to do this and you can fine tune it afterwards. Don’ t be afraid to zoom in to 1:1 to find a nice white spot.

You can also tweak it by hand. Open your color contrast and you will see “green vs. magenta” and “blue vs. green”. Right-click on the sliders to get the mouse adjustment curves or inch your sliders left and right very carefully. You should be able to get the same results as with the white balance. When adjusting a landscape photo, keep your eyes on the clouds, Pink clouds are not uncommon, but you do not see them in the middle of the day.

Don’t be afraid to open the same picture in another editor like RawTerapee and adjust it there. Different tools have different algorithms and sometimes one makes it perfect in one go. You can now mimic those settings in Darktable to get your perfect picture. Once you have your white balance sorted, you can

start layering things on like HDR and masks. However, keep your eye on your histogram always.



That's the low down-and-dirty on white balance. Obviously there is a lot of theory and practice behind it, but this should be enough to get you up and running without bogging you down in the details. It is instant results that will speed up your workflow. When the workflow is a slog, nobody enjoys it. Darktable post processing is actually relaxing therapy. Enjoy it.



Questions? Comments? E-mail us on misc@fullcirclemagazine.org

Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.





[RETURN TO CONTENTS](#)

How To - Drawing with Inkscape - Part 93

Written by Mark Crutch


Last time, we began work on an SVG file which uses JavaScript to animate the “transform” attribute of an object when loaded in a web browser. We’d got as far as animating the rotate() and skewX() functions, and you’d been left with the challenge of adding skewY() to the mix. Hopefully you worked out that this was mostly a case of copying the existing code, and replacing “X” with “Y”. Specifically:

- *Create three new properties in the animProperties object: skewYDuration, skewYMin, skewYMax.*
- *Insert another block of code to calculate the current value of the skewY() function, given the supplied timestamp. This is literally a copy of the skewX code with the letter replaced throughout, resulting in a skewYAmount variable at the end.*
- *Add another line to the setAttribute() call to include the skewY() function in the transform attribute, passing the value of the skewYAmount variable via the template string.*

I also left you with something of a puzzle: with the addition of the skewX() function, the square doesn’t just rotate and skew in the middle of the screen as you might expect; instead it moves wildly in and out of the browser window as it rotates. Adding the skewY() function simply exacerbates the problem. Why does it do that? And how can we get the behaviour we expected?

The reason is quite simple: whereas the rotate() function has an optional pair of parameters for setting the center of rotation, there’s no equivalent for the skew functions. Skewing takes place relative to a baseline, rather than a single point, but there’s no generic skew() function that lets you specify this baseline via two sets of coordinates. Instead, there are only two

possible baselines available: the x-axis (via the `skewX()` function) and the y-axis (via the `skewY()` function). The graph below shows the effect of skewing along the x-axis for the same size of object positioned in a variety of places:



The red squares along the x-axis all skew “in-place”, resulting in the outline shapes displayed. The green and blue squares – colored to avoid confusion where they overlap – move to the left and right as a result of the skewing process. It doesn’t take much thought to realise that the amount of movement to the left and right is actually proportional to the distance from the x-axis, so even a small angle of skew can quickly move a shape by a large distance if it’s located far from the axis. And, of course, these rules also apply for the `skewY()` function, but rotated by 90° .

You may recall that we did some deliberate manipulation of our object in order to position it in the middle of the screen (50, 50). Unfortunately, all that work means that skewing the shape also pushes it around. The problem is that we have two conflicting requirements:


1. We want the object centered at (50, 50) for display.
2. We need the object to be centered at (0, 0) in order to skew it.

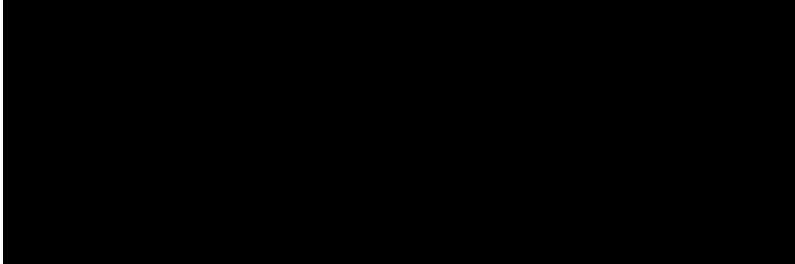
There are a few ways to solve this conundrum:

- *Use a `translate()` function to move the object to (0, 0); then skew it; then use another `translate()` function to move it back again.*
- *Change the `x` and `y` attributes of the object so that it starts at position (0, 0). After skewing, add a `translate()` function to move it to (50, 50).*
- *Use a `matrix()` function rather than the `skewX()/skewY()` functions, as this can combine skewing and translating into a single call.*

I'm going to immediately dismiss the last option, as it requires far too much mathematics. But it does, perhaps, explain why Inkscape always uses `matrix()` rather than the named skew functions. The first option is probably the clearest in terms of what's happening, but it results in the position of the object being calculated three times: once with the initial `x` and `y` values set to 50; once when it's translated back to (0, 0) prior to the skew; then a final time when an inverse translation is applied to put the skewed version back into place.

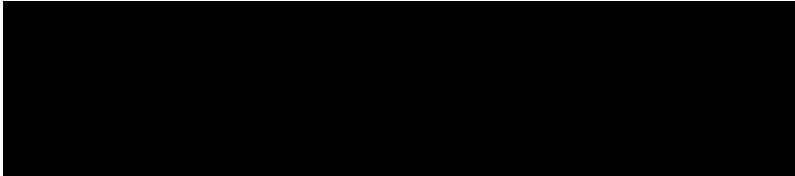
The best compromise for this project is the second approach. Simply changing the `x` and `y` values in the SVG, however, means that the source file no longer holds the “true” values for the coordinates, so if the JavaScript fails to load for some reason the square will be positioned at the top left, rather than in the middle of the screen. A slight modification to this approach – and the one we'll take here – is to leave the SVG file untouched, but change the `x` and `y` values to zero at the start of the JS file. That way, if the script fails to load you still get a stationary fallback image with the square in the right place, but if everything loads correctly, the JS immediately changes the object's coordinates to make for less work in the rest of the code. The first step is therefore to modify the start of the `initialise()` function to get a handle to the `<rect>` inside the group, and reset the attributes:



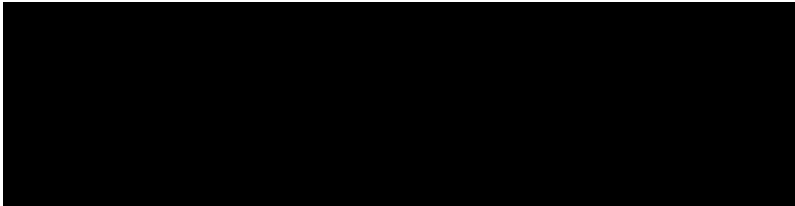


We use `group.querySelector()` to get the rectangle as it limits the search to descendants of the group object (compared with `document.querySelector()` which would search the whole document), and it makes it easy to replace the argument with an ID or class selector when working with a more complex drawing, or with a different element selector if we change the content of the group to be a different type of object (e.g. a `<path>` or an `<image>` instead of the `<rect>`). Once we've got a handle to the element, it's back to our old friend `setAttribute()` to set the new values we want.


Reloading the page at this point shows that, if anything, we've made the problem worse! Now the square is rotating over an even wider range than before, spending most of its time out of bounds. The reason for this is also a simple one: remember that we used the three parameter version of the `rotate()` transformation function, so even though our object is centered at (0, 0) it's still rotating around (50, 50). Now we can pare that function down to the single value version, and the line where we set the value of the transform attribute becomes this:



Reload the page and we've got what we expected: a square that rotates and skews in-place at the top-left of the screen. Our last step is to move it back to (50, 50) with an extra `translate()` step inserted to the start of the list, whose values are hard-coded:



Note that the transformations are actually applied in reverse order: first the skewY(), then skewX(), then rotate() then finally translate(). When all we had was rotate and skew functions the order made little difference, but adding the translate() makes a huge difference. Put it at the end of the list, and we're back to the same problem with the square zooming around as it skews. At the start of the list, however, we've got a nicely controlled square, rotating and skewing whilst never leaving the middle of the screen.




To complete our set of transform functions, let's add a scale(), so that our square also grows and shrinks. This function can take one or two parameters to indicate the scale factor: if only

one is provided then the object is scaled equally in both the x and y directions; if you wish to scale the two directions differently, then you have to provide two parameters.


Note, however, that there's no parameter for the center point of the scaling operation. As with the skew functions, your object has to be positioned with its center at (0, 0) if you don't want it to move as well as change size. Since we've already handled this problem for skewing, we just need to ensure that our `scale()` function is put into the transform attribute after the `translate()`, to ensure that the scale operation is performed before the translation.

To make our animation more interesting we're going to animate the change in x and y scale separately, over different time periods but within the same range of 0.1 (one tenth of the width or height) to 3.0 (treble the width or height). We'll encompass these parameters as two more sets of properties in the `group.animProperties` object:




Like the skew functions we also want to animate from the minimum to the maximum, then back again – as opposed to continuously going in one direction as we did with `rotate()`. We therefore need a couple of blocks of code to calculate the relevant value at any given time point, changing direction after each period. Here's the code for the x-axis scaling – compare it to the equivalent block for `skewX()` from last month and you


should be able to reproduce code for the y-axis scaling yourself.



And, of course, we need to add our `scale()` function and two new parameters to the transform attribute:




At last we have our object rotating, skewing and scaling, all while centered in the browser window – though a static screenshot doesn't really do it justice.



I'm going to finish this month with a couple of exercises for you to try, which build on the animation we've created over these previous few articles:

1. *Our final transform attribute has a fixed `translate()` function to position the square in the middle of the screen. Why not add another two sets of parameters to also animate the `x` and `y` position, causing the square to move around the window a little as well. Setting min and max values either side of 50 means you can replace the hard-coded coordinates in the existing `translate()`. Or you could have a negative min and positive max, then use the values in a second `translate()` function – but be careful of the ordering!*
2. *Try replacing the contents of the group with something else. It could be a more interesting single shape, such as a star or more faceted polygon, but it could also be any other Inkscape*

drawing – with multiple shapes and colors. Simply replacing the `<rect>` with an `<image>` element makes for an interesting effect, reminiscent of the kind of thing that required a Hollywood budget back in the 1980s.



The most important thing to remember is that this animation code – and the JS that we used previously in this series – are just examples to get you going. There's no reason why you can't create an animation that messes with the transform attribute whilst at the same time altering the fill and stroke, or manipulating the "d" attribute of a `<path>` in order to change the shape being drawn. With an understanding of how to change attributes and properties from JS you can create interactive or animated SVG images that go way beyond the frame-based limitations of an animated GIF. It's a bit of a cliché, but the only real limit is your own imagination.

Mark uses Inkscape to create three webcomics, 'The Greys', 'Monsters, Inked' and 'Elvie', which can all be found at <http://www.peppertop.com/>

RETURN TO CONTENTS

Using Krita to Rework Old Photos - Part 3

Written by Alan Ward


This series is aimed at learning to make something of the old photos in my possession, and others in the public domain due to their age. You, the reader, are welcome to tag along, and, I hope, glean some small insight and perhaps an idea or two from time to time. No promises are made as to quality of the content, or potential errors and omissions. I am a computer scientist, not a true artist or a professional of image restoration. So please take all this as a best effort, but with no firm guarantees — much as is the case of most open-source software.

In this part of this series, we will continue work on a simple landscape, a photo of the castle of Foix in southern France about the turn of the century. With the passage of time, this photo is now stated to fall within the public domain, has already been digitized by the Rosalis project of the municipal public library of Toulouse, and may be downloaded from Wikicommons at address: [https://commons.m.wikimedia.org/wiki/File:Ch%C3%A2teau_en_ruines_\(8056081904\).jpg](https://commons.m.wikimedia.org/wiki/File:Ch%C3%A2teau_en_ruines_(8056081904).jpg)

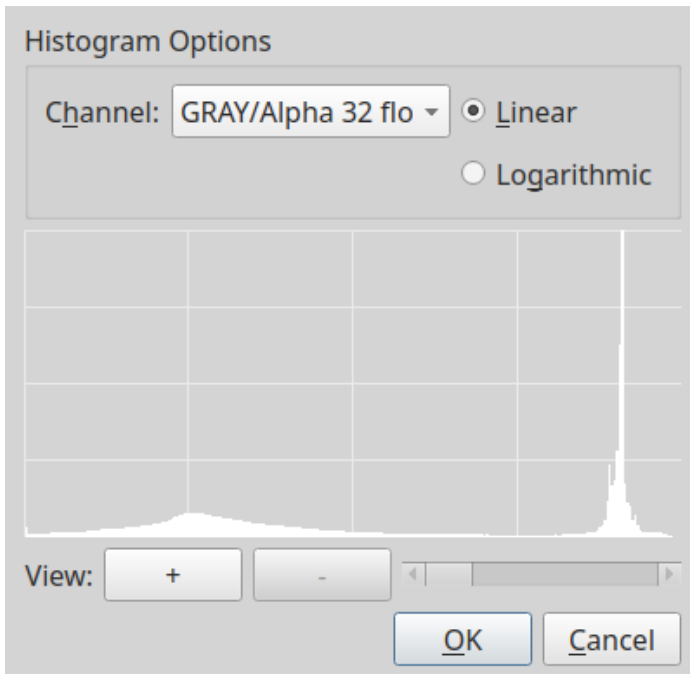
In the previous episode, we cropped the image slightly to remove some damage along the lower and left edges, used the airbrush to clean up more damage near the top edge and some typical dark spots that were quite apparent on the clear sky, and then an overlay layer along with a dark color and the airbrush to reduce the effect of a large whitish blemish in the lower left corner.



This time, we will try to solve the issue of the dark castle and trees to the right of the photo, and put some color back into the scene. Before proceeding, we will make a copy of the Krita file, to save the current state for future reference or as a means of recuperation if disaster strikes. In our new copy, we will take the opportunity to combine all three layers present (the original image, a normal color layer and an overlay), to work on a single item going forward. Simply do a right-click on the topmost layer, and choose “Flatten Image”.



It is always instructive to take a quick histogram of our layer. With our only layer active, choose menu option “Layer”, then “Histogram”. In the resulting plot, we may need to change our channel. In my case, I needed to choose “GRAY/Alpha 32 float” since the original JPEG file was in grey-scale coded as 32-bits. The appropriate setting will be immediately apparent once the curve appears beneath.



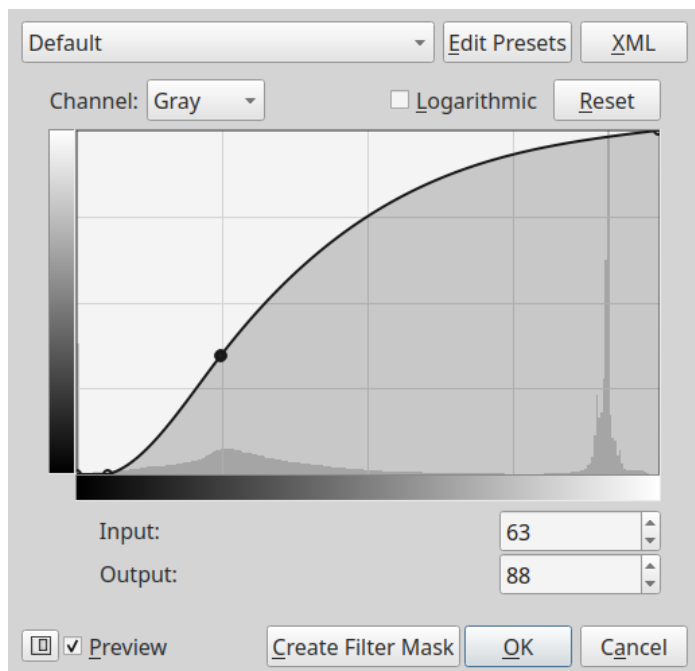
In this plot, the amount of pixels in each color is plotted, with frequencies of dark colors to the left, and light colors to the right. For our picture, we can observe a sharp peak to the right, which corresponds to the large amount of light-colored pixels in our sky, that cover around a third of the total image area. Towards the left of the histogram, the frequencies of darker pixels are more stretched out, denoting a scale of several dark colors that go from pure black at the extreme left, to a middle gray at the center of the plot.

In other words, the darker part of our image –the landscape itself– is quite well-balanced, with pixels of many shades of gray. However, sky colors are very far apart from the earth, with a clear separation of a region of lighter colors in which there are not very many details. This results in a high level of contrast. For a more pleasing result to the eye, an ideal situation would be that in which the histogram plot would be as flat as possible, with a wide range of color values giving a richness of details.

Moving forward, we have one basic choice to make. Do we leave the sky as it is, in sharp contrast to the ground, or do we try to compress sky colors into a narrower band, thus freeing up some color space that we can use to spread out earth colors somewhat. In our case, our sky is already quite washed out with little detail left in our original photo, so I feel that by compressing it further into even fewer colors we are actually losing very little detail.

To alter the color balance of a layer in Krita, go the menu option “Filter”, then “Adjust”, then choose –for instance– “Color adjustment curves”. As can be seen in this submenu, the curves are not the only option available, and each user may want to give other tools a quick glance.

In the curves’ dialog, the diagram in the middle represents the transformation we will apply to colors. The horizontal scale represents current colors, and the vertical the colors into which they will be transformed.




The curve itself is initially a diagonal line. This means that, at

first, each current color will be transformed into itself — i.e., it will be left as it is. Now, by simply moving the curve (and so introducing control points represented by dots), we will tell the application which alterations to perform. Moving the curve upwards will make the affected range of colors lighter, and moving it downwards will darken them. In our case, the general intent was to lighten the earth, so I raised the curve lightly. It is best to avoid wide, sweeping movements: regions where the curve approached being either completely vertical or completely horizontal tend to give rather strange results that shock the eye. Please note, also, that I have retained a second control point low in the left corner: this is here to ensure that pixels that are completely black remain so.

Our resulting picture, so far is:



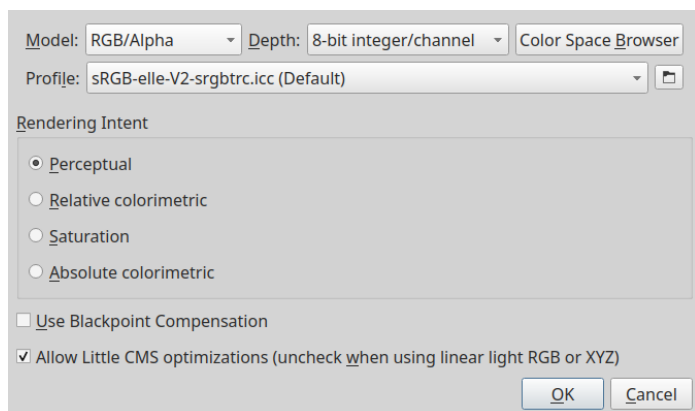
If we consult our new histogram, we can see that the main image has colors spread out over a wide range, while the sky is still bunched up in a single sharp peak to the right. This last bit does not represent a major alteration over the original state, while the main body of the image is more balanced than previously.



So far, our photo has progressed quite a bit in quality in comparison to the original. However, one aspect is lacking to give the scene a tad more realism: colors. These were lost during the initial photography, when the black-and-white negative was exposed. In a broad sense, the chemicals spread over the glass plate were sensitive only to the general light intensity falling on them. Information was lost concerning the exact wavelengths. In more recent color photographic film, several layers of sensitive chemicals inter-spaced with filters allow retaining information on separate colors on, at least, three different wavelengths. Thus, a (rather good) approximation of the scene's original colors may be recomposed after processing, making positive copies and development. In our case, however, this color is completely lost to us. In digital terms, what we now see as a level of gray coded 200 on a range of 0 (black) to 255 (white) shows up as a light gray. It is possible that the original color in the scene was a light gray, with red, green and blue values of 200, 200 and 200 and with average intensity 200. But the original color could very well have been 250, 200 and 150, a light salmon color due to the presence of more red and less blue. It could also have been 150, 200 and 250, a light pastel blue, or any other variation of primary colors giving the same average intensity of 200.

Our takeaway here, is that in order to colorize our image, we will need to add back the information that has been lost. Luckily, our experience as humans tells us that the sky should probably be blue, and the trees green.

To finish up for today, let us add in Krita a new transparent layer. We will also need to convert the image from gray-scale to color. To do so, choose menu option “Image”, and “Convert image color space”. Set it to “RGB/Alpha”, i.e. three color channels (red, green and blue), plus an alpha (transparency) channel.



Now, with the airbrush at a low opacity setting (about 0.2 should give good results), and rather broad opening (perhaps in the range of 40 to 300 pixels, depending on overall image size), let us brush in some color. Unlike in the first part of this series, when we were covering up small defects, now we will need to proceed with wide open strokes giving coverage to large areas at once, but with little density. For our first attempt, we will be using just three colors –a light blue, a darker green, and a reddish brick pink– and not worrying about complete coverage of each part of the scene.



The result, so far, is quite reminiscent of period hand-coloring of black and white photographs in the early 1900s. We already get a better appreciation of how the city's rooftops and white buildings must have stood out from the surrounding vegetation and blue sky. However, we can also see that this approach has some drawbacks. The main one is that, since we are simply adding color on top of an existing black motive, we are darkening considerably the scene as a whole. This is not very realistic. Another problem is that our green coverage, for example, is uniform across different types of trees and grass, each of which would have had its very own shade and intensity. So, we need to think about how to modulate our new colors as a function of what they are laid out upon. This is what we shall do in the next episode of our series, though the reader is encouraged to experiment freely on his/her own beforehand. Until then, take care!

Alan holds a PhD in Information and the Knowledge Society. He teaches engineering at Escola Andorrana de Batxillerat (high-school). He has previously given GNU/Linux courses at the University of Andorra and taught GNU/Linux systems administration at the Open University of

Catalonia (UOC).

[RETURN TO CONTENTS](#)

THE DAILY WADDLE



The Daily Waddle

They are going to make
a movie here!!

How do you know?

There's VLC icons
all over the road!



[RETURN TO CONTENTS](#)

Linux Loopback

Written by S. J. Webb

I intended to write further about the Lumina Desktop. However Project Trident recently decided to leave the FreeBSD base for its OS. They will continue to operate using Void Linux for further development. The TrueOS is no longer a workstation-focused product as being developed by iXsystems, TrueOS is now a server OS almost. Consequently, it is hard for Project Trident to continue using the FreeBSD base.


They cite Void Linux being the most BSD like Linux, since a former NetBSD developer created Void Linux. The Project Trident developers hope a synergy relationship will develop between Void Linux and their work.

I hoped to continue to use Project Trident for the ZFS file system, Lumina Desktop, and the BSD kernel obscurity. Yet I wonder if ZFS will work on the Linux kernel. I understand that Lumina works flawlessly with BSD, but will it be marginally operating on the Linux environment? In the end, I want to use BSD for one of my laptops.

So the desire to use BSD is now making me look into other options. I could use OpenBSD, GhostBSD, or MidnightBSD. I never ventured into the realm of OpenBSD. It is the second most popular BSD OS. It is known for being rock solid, reliable, and easy to install. Additionally, it is focused on being secure. However FreeBSD and OpenBSD require a configuration file skillset that most end users lack. So I am left with fewer options.

I like GhostBSD and MidnightBSD, but both are the easy BSD options. There is a new OS called FuryBSD. FuryBSD is a new BSD on the scene. The developer, Joe Maloney, is an iX systems developer. It is only a few months old. It does not rely on the TrueOS platform, but relies solely on FreeBSD. The default DE

is XFCE.



I installed FuryBSD without issue onto the laptop. I will have to configure my WiFi Card for the laptop. I hope to follow the instructions from the FreeBSD manual for activating my WiFi. I did not notice there is no software center for FuryBSD. I will have to use the terminal to install packages, system updates, and software apps, to make the OS fully functional. Further details can be found here: <https://www.furybsd.org/>

S. J. Webb is a researcher coordinator. When he is not working, he enjoys time with his wife and kids. He thanks Mike Ferarri for his mentorship

RETURN TO CONTENTS

Everyday Ubuntu - Personal Information Management Part 1

Written by Richard Adams

Hope everyone is having a great new year! Now seems like a very apt time to talk about personal information management in the light of a very common New Year's Resolution, the commitment to become better organized in 2020.

Program Choices

Personal Information Management, being a core part of computer functionality for most computer owners and users, it should come as no surprise that there are multiple Personal Information Manager applications available on Linux. Common and well-regarded ones include Evolution, Kontact, and Lightning/Thunderbird.

I was always partial to Kontact back in my pre-Ubuntu Linux days when I was running SUSE or Fedora. Coming from a Windows and Microsoft Office background, I found Kontact's interface similar to Microsoft Outlook and therefore comfortably familiar. However, upon switching to Ubuntu, I found that Thunderbird was easier to set up and not very hard to learn, so I switched to it.

An important aspect to most Personal Information Managers, or PIM's, is their integration with email. Of course, Thunderbird is primarily thought of as an email client, but it does have some PIM capability built in 'out-of-the-box', and gains even more if you install the Lightning plug-in. I have to chuckle these days at the memory of using Schedule+ as my PIM back in 1995 or so, and being singularly unenthused at the upcoming Outlook 97, as I saw little or no connection between PIM functionality

and email. Live and learn, I guess.

Thunderbird also has the advantage of a solid cross-platform footprint, meaning that once you learn it on Linux, you can carry that knowledge and familiarity over if you switch to MacOS or Windows at some later date, or if you use multiple operating systems on a regular basis, like I do. So, due to simplicity and universality, we're going to go with Thunderbird as our primary PIM and email client, at least for the time being.

Thunderbird Installation

Thunderbird can be installed using Synaptic Package Manager or apt-get, the command-line installation tool. For details on how to install software, see Everyday Ubuntu in Full Circle Magazine #130. We're going to use apt-get here, because it's very easy and straightforward. It is a command-line, or text-based, tool, but it's really easy to use, so don't be afraid!

Go to the terminal (usually the third icon from the top on the Launcher that runs down the left-hand side of the screen), or go to the Dash – top icon on the Launcher on screen left:



and type in term, then click the terminal icon:



and from the terminal type:



then hit the Enter key on your keyboard. The command `sudo` stands for super user do, and tells Linux that you want to perform a command that's allowed only to administrators, or superusers. Enter your Super User or administrator password that you set up during Linux's initial installation. This will start a Super User installation routine for Thunderbird. You will need to be connected to the internet for this command to operate, as it will have to download Thunderbird's installation files. Once it's complete, we're ready to set up email.

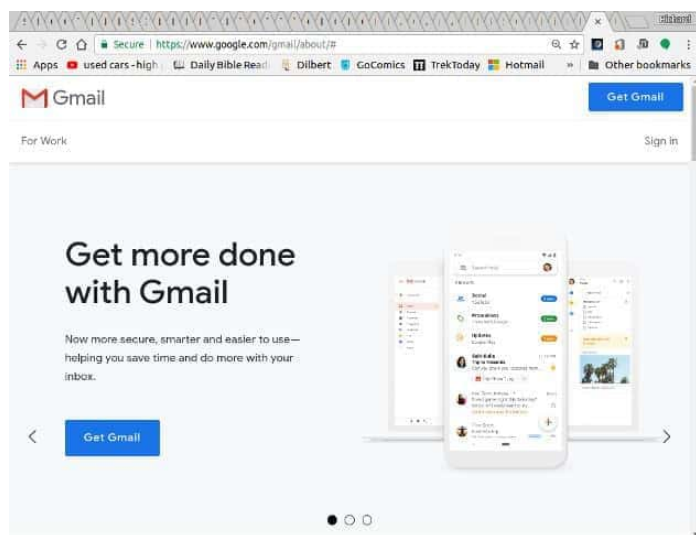
Email Setup

Now we're ready to set up an email account in Thunderbird. Your Internet Service Provider (ISP) may have given you an email account, in which case you'll have to get setup information from them. If you don't already have an email account, or want to set one up specific to your machine, you can get a free email account online from Google via their Gmail service. It's free, and it's a flexible service. Microsoft has the similar Outlook/Hotmail infrastructure, but it has only recently added support for POP protocol directly. POP stands for Post Office Protocol and is the protocol for email that we shall be using here, although a protocol called IMAP (Internet Message

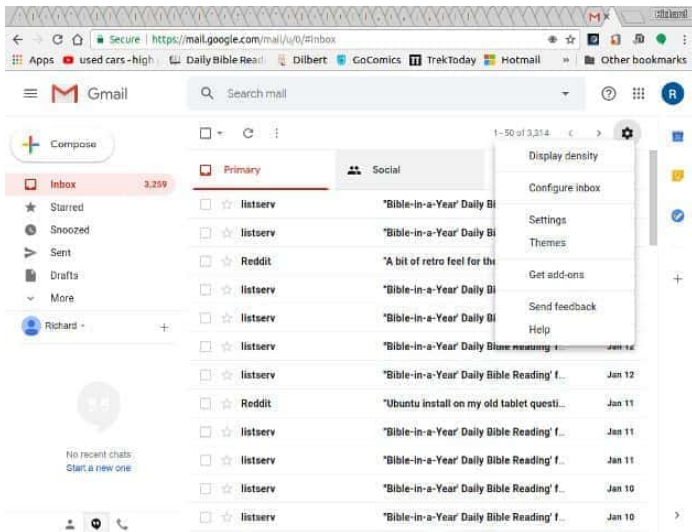
Access Protocol) would also suit our purposes almost as well. The idea of selecting one of these protocols is so that we can download the emails to our local email client, Thunderbird, so we don't have to be online to read them, and POP is a better choice for that. Thunderbird will download your emails when you are connected to the internet, then you can read them later whether you are still on the internet or not. Gmail also has the advantage of accessing the robust Google infrastructure so you can readily share information across different systems and devices.

Gmail Setup

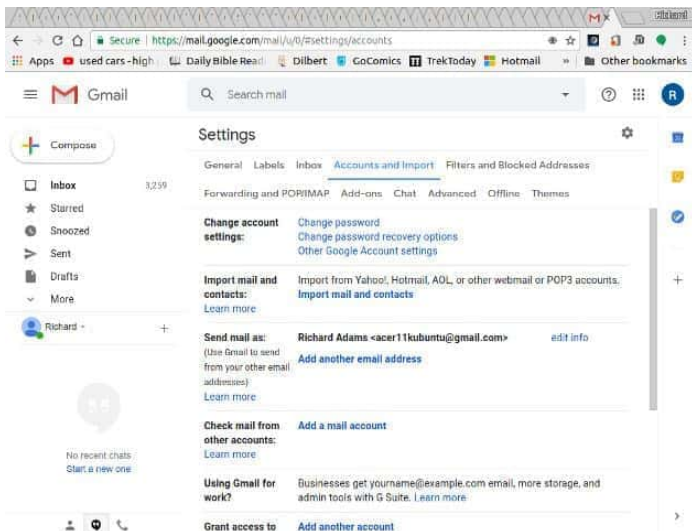
First, fire up your web browser of choice, and go to [gmail.com](https://www.google.com/gmail/about/#). You should see a screen allowing you to sign in or sign up.



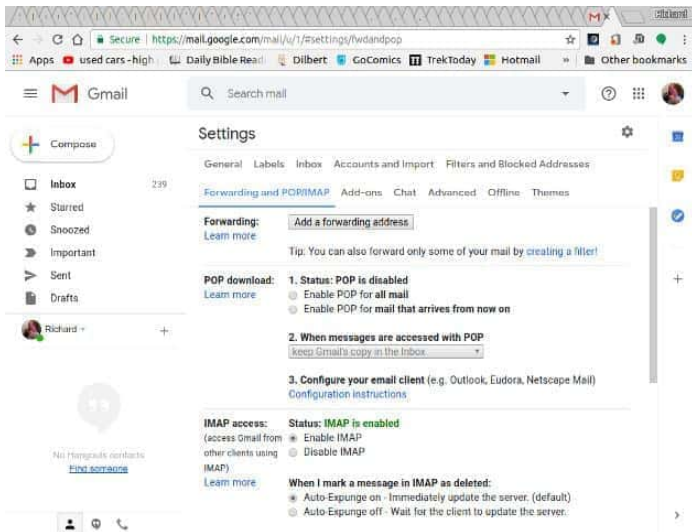
Select the option to sign up and fill in all the information required. Once this is complete, you can sign in to Gmail in your browser and set up POP mail capability. To do so, go to the gear icon near the top right side of the screen and click it, then click Settings:



On the Settings screen, click the tab for Accounts and Import, then Forwarding and POP/IMAP:



In the POP download section, choose Enable POP for all mail:



For number 2, select whether or not you want email to stay in Gmail's inbox (a good choice if you're going to use the account on multiple devices or access it via webmail/browser), or to delete from your online mailbox after you download to Thunderbird, or to archive from the online inbox after accessing via Thunderbird. Ignore the 'mark as read' option for now.

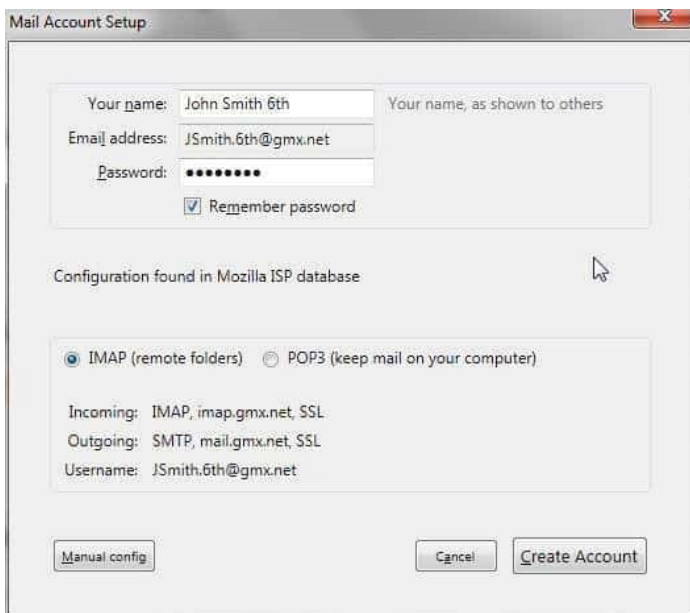
Thunderbird Client Setup

Now we're ready to set up Thunderbird to access our new Gmail account. Use the Dash again to look up Thunderbird and click it to launch. On first run, Thunderbird will ask you for your account information.

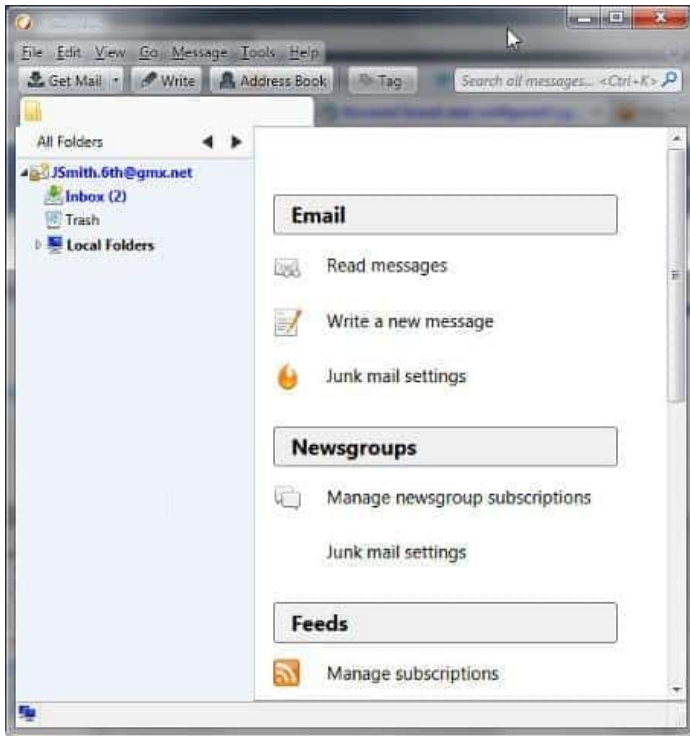
Hit Continue after filling out the required info. You'll get this dialogue box;



Click POP3, then Create Account. As you can see, the Thunderbird servers already know configuration information for various ISP's and email providers, including Gmail, so you can just provide your Gmail (or other email account) information and Thunderbird will set up everything for you – no muss, no fuss. Nice!

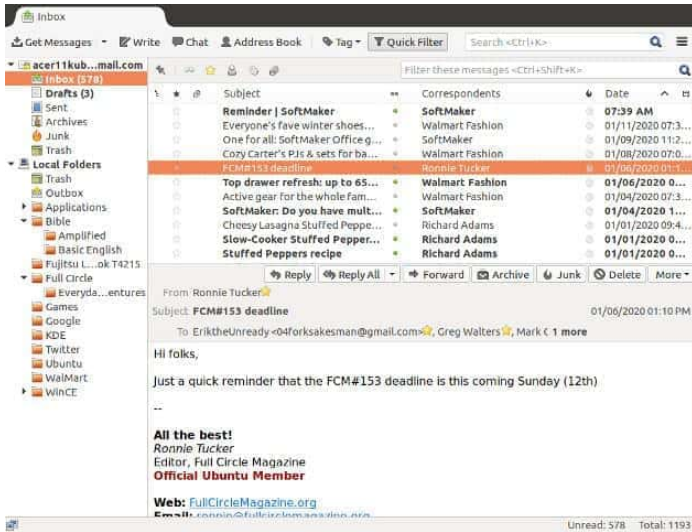


Thunderbird will now show you its main window:



Click Read Messages to see your Gmail inbox;

And that's it! Now your email is set up to be accessed via Thunderbird! Next month, we'll cover how to go beyond email to manage your personal information in Thunderbird.



Next month: Setting up Lightning to extend Thunderbird's PIM capabilities.



Richard 'Flash' Adams lives in rural north Alabama and has saved the world from the Kilrathi, the imps and cacodemons of Doom, and Shodan of Citadel Station. You're quite welcome. He used to enjoy Cowboys and Falcons NFL football, but this season? Oy.

RETURN TO CONTENTS

Ubports Devices

Written by UBports Team

Coming soon

[RETURN TO CONTENTS](#)

THE DAILY WADDLE



The Daily Waddle

True ! False ! True ! False !

Stop boolean me!



[RETURN TO CONTENTS](#)

My Opinion - Container Security Concerns

Written by Chris Binnie

There are a number of areas to keep in mind when it comes to securing containers that you have running on your Ubuntu hosts and their workloads. In this article, we'll have a look at a high-level overview of the key things to bear in mind. We'll start with a more traditional security approach first and then focus on some of the other pertinent aspects.

Host Security

If you're completely embroiled with getting your containers running just as you'd like them to, then it's easy to overlook the fact that you should firstly pay close attention to the host machine that those containers are running on. Clearly the security of the host that your containers are relying upon determines how much uptime you will achieve for your applications and additionally affect the number of threats that might cause issues for you.

The usual rules apply when it comes to your host machines. You will need to update packages when they become available, most likely schedule reboots for kernel updates and also ensure that a minimal number of packages are installed to help limit the attack surface.

Aside from the time-honoured package updates, you should ensure that only specific network ports are exposed, locally and publicly, and then firewall off all other points of network access. Where possible, rate limiting access to applications, and monitoring accesses via public network ports in detail, is also advised.

Limiting Crosstalk

Another area can be helped significantly with a smattering of common sense and logic.

Imagine for example that you have three containers running on a single host, each with a unique application that provides a service of some sort. It's advisable to carefully consider how those containers might interact from an architectural standpoint. If Container A passes data to only Container B, then Container C doesn't need access to Container A directly at all and should be isolated at both a host and an internal networking level.

Consider another scenario where you might have two upfront web servers running via two containers, and additionally a single backend database server. The front-end servers pass read and write requests back to the database and only TCP port 443 needs to be exposed to the public from a network perspective.

There are however lots of other containers running on the same host machine which shouldn't have any visibility of the potentially sensitive database traffic being passed through the database server before being written and stored outside of the container. The recommendation here would be to use a bridge network. By connecting only our two web servers and single database server to that private network, we will successfully limit network access and provide a layer of isolation.

This means that if another container on the host is attacked and compromised there's more layers of security potentially for the attacker to break through in order to get access to the database data.

Common Vulnerabilities

The dreaded CVEs (Common Vulnerabilities and Exploits) apply to package updates just as we saw on a host's Operating System a moment ago and you should monitor CVEs using any of a variety of tools.

How you approach updating packages inside your containers is a different matter. There are a number of valid approaches but

a strategy is definitely required and often dismissed as being trivial. As a result, ad hoc patching practices can become dangerously erratic. This is because even the very popular container images hold a surprising number of CVEs at the best of times. Indeed many packages don't have fixes available from the vendor at the time that a container is spun up and potentially exposed publicly meaning patching is near to impossible.

The recommendation is to think carefully about what risks affect you the most. Take time to fully understand your priorities and the attack surfaces that you present inwardly and externally. Then think about how realistic updating every package for triggered alerts is and at what frequency. Choosing a tool which can automatically alert you to software updates can be a little daunting as there are a few. One such sophisticated tool which can run from containers itself can be found here (<https://anchore.com/opensource/>) and is called Anchore.

As mentioned, not all issues with CVEs can be mitigated against using package manager updates because fixes haven't been released by vendors. Trade-offs between using alternative base Operating Systems and using alternative applications might be required on occasion.

Locking Down Access

As all containers share the same host kernel on a machine, it's imperative that suitable isolation is put in place to protect the host. Without the host functioning correctly, ultimately there's only downtime which means failure to stop a successful attack on one container from performing a "container escape" will mean that other containers, and their applications, or even the host itself, might succumb to the attack.

The Linux kernel has introduced a number of clever isolation techniques over the years which we won't look at in detail here.

You should be aware of Kernel Namespaces, however, which

means Customer A can't see what Customer B is doing if all goes well from a security perspective. Additionally, you should know about Control Groups, known as "cgroups", which can apply strict resource quotas such as how much memory or disk access a container might be allowed.

You might want to also read up on Kernel Capabilities – which can achieve all sorts of useful access restrictions. For example, do you really want a container to be able to change the time and date on your host's system clock? It's unlikely!

Although we've just scratched the surface on this topic, you should also make sure that lesser privileged users can't start and stop containers with permissions that an attacker can do bad things with. Leave that sort of thing to the "root" superuser alone.

Orchestration

When you're running more than a few containers at once, it can be a little like herding cats trying to get them all to behave properly.

Many people turn to Kubernetes (<https://kubernetes.io>) for larger workloads, and this brings its own security considerations in addition.

In brief you should use a more sophisticated approach to network isolation than we looked at earlier, and ensure customers or applications are split up by namespaces.

You should also finely tune Pod Security Policies, which are cluster-wide, to limit crosstalk noise and additionally host access.

Finally for sensitive container scenarios, you can introduce Security Context Constraints (<https://kubernetes.io/docs/tasks/configure-pod-container/security-context/>) on a per Pod (one or more containers) basis.

You're encouraged to read further in order to get cluster

security working well, amongst a number of other areas.

Stronger Isolation

Despite a very high level of isolation being possible through a previous Virtual Machine incarnation of “rkt” (<https://coreos.com/rkt/>), the container runtime by CoreOS using KVM (https://www.linux-kvm.org/page/Main_Page), all the rage currently is Kata Containers (<https://katacontainers.io/>). Essentially, Kata Containers’ aim is to provide Virtual Machine levels of security protection (where true host and workload isolation takes place), and allowing containers to benefit from such protection.

Virtual Machines adopt a hardware level of isolation making attacks much, much harder than just circumventing a host machine’s kernel. By cleverly enforcing that level of isolation, but enjoying the quick start-up times, portability, and predictability of disposable containers, it is certainly a technology to pay close attention to. Stay tuned.

The End Is Nigh

We’ve barely scratched the surface in terms of getting into the detail about securing applications in your containers.

Hopefully, however, the key areas which we’ve looked at briefly will give you some food for thought about what to read up on further the next time you need to make a decision about how to approach solving a container security problem.

Chris Binnie’s latest book, ‘Linux Server Security: Hack and Defend’, shows you how to make your servers invisible and perform a variety of attacks. You can find out more about DevSecOps, containers and Linux security on his website:
<https://www.devsecops.cc>

[RETURN TO CONTENTS](#)

My Opinion - Linux on the Go

Written by Knightwise

My name is Knightwise and I own a Surface Go. There, I said it. I'm part of the rather controversial club of people who forked over 600+ euro's of cash on Microsoft's underpowered entry level gateway drug to the Surface Line. A device that is either loved for its versatility, its form factor and its overall cuteness by some, and hated because its an over promising under-delivering piece of 'garbaggio' by others.

Me? I'm still on the fence. Having owned the device for over a year now, I have seen it go through cycles of daily heavy use, and weeks of neglect. Sometimes I'm in one camp, sometimes I seem to be in the other. A couple of weeks ago I decided on a different approach. Since I already own a pretty powerful Windows Machine (I have a Lenovo X1 Carbon as my daily driver), maybe I should take a look at running Linux on this little machine (the Surface Go) and see what happens.

The hardware compatibility of the device allows you to run almost any Linux distro on this machine – out of the box. I say almost because there are a couple of things that aren't working. Wifi is one of them and it's important to know this before you “nuke and pave” your hard drive. The fix is easy, download some files and copy them over to some directory, and boom, bob is your uncle. But make sure to do it either before you install Linux or do it on a separate computer. The Surface Go has no classic “ethernet port”, so popping in a network cable is not that easy.

Everything else just works or is provided by installing a custom kernel written by a great guy named Jake. I could lie and say it's hard to do, but it's not. Jake has written a simple script that you copy and paste into your terminal and boom: Screen

rotation, stylus support and all the other goodies, but one, just work.

Hardware Support

The one thing that does not work yet are the webcams. And that is a darn shame because the webcam quality on the Surface Go is downright awesome. Seriously. The cam quality on this little laptop puts the one on my Macbook pro to shame at 1/3rd of the price. I know the programming geeks who reverse engineer the Linux drivers on this thing are probably recluse and shy, so a webcam is not as high on their list as it is on mine, but I would like to see this fixed pretty soon.

Screen

The only additional tweak I would recommend is setting the screen scaling to 150% (and you need a gnome command for that). It's not that hard to find, but worth it since a 100% scaling makes the text on this full HD 10" display too tiny, while the 200% makes it too large.

Performance on this baby under Linux is even better than under Windows. Some CPU heavy processes (like permanent file indexation) don't run under Linux, so the device is pretty snappy. Battery life is not as optimal as under Windows, but still gives you a pretty fancy mobile Linux machine. There is the possibility to play around with suspend/hibernate when you close the lid, but Ubuntu's sore thumb (no proper support for hibernate) still sticks out here. The device boots lightning fast, so having to 'power down' to save battery is not the worst thing in the world.

I went for the Surface Go with the 8 GB of RAM and 128 gigs of hard-drive space. The linux distro does not eat up much space so there is plenty of room for your data. I went for the "minimal Ubuntu Install" and "picked and chose" my favorite mix of Command-Line and Gui Apps. Additional storage was added by jamming in a 128 gig micro SD card in the back and I have successfully doubled my storage capacity.

Ports are pretty scarce on this thing. One USB C port (and the well known ‘Surface’ connector) is about all you get. You can even charge the Go via USB and hook it up to an external USB C hub (or a dock). When jamming it into my USB-C HP Dock, I can get it to power one 25-inch display as a “second screen”, but powering 2 x 25-inch displays is a little beyond the capabilities of the Go.

The one thing that makes the GO pretty sweet is its size. It’s slightly larger than a regular iPad (and has a very unique form factor), but smaller than its 12-inch Surface Pro brothers. This is a very important factor for me when using it as a touch device. I flip the keyboard back and I have a light and comfortable device to read, watch videos, and scribble down some notes in Journal + + The weight and the balance of the device makes it a pretty cool ultraportable lappie.

I’m not a developer so I can’t say anything about compiling speeds and such, but I DO work in the “corporate world” as an IT consultant. I am pretty sure that if they fixed the webcams I would have a sweet machine to survive out there. I can run an entire Office365 environment in my browser, Install the Citrix Client, and no-one of my suit-and-tie-wearing friends would be any the wiser that my Go is actually dancing with the Penguin. (I don’t wear a suit and tie by the way :p).

In the end, I mostly use my Surface Go (under Linux) as my little “drag along” Linux laptop. Gnome 3 is not a bad OS for touch-enabled devices, but I do tend to use the keyboard and mouse most of the time. Linux on the Surface Go surely turns the little tablet/laptop into a very interesting device that, like any other touch device without a “real” mobile operating system, sits somewhere in between awesome, quirky and useless, depending on your mood.

Links :

- *The SurfaceLinux Subreddit:* <https://www.reddit.com/r/SurfaceLinux/>
- *Jakeday’s Kernel page* <https://github.com/jakeday/linux-surface>

[RETURN TO CONTENTS](#)

How To - Write For Full Circle Magazine

Written by Ronnie Tucker

Guidelines

The single rule for an article is that it must somehow be linked to Ubuntu or one of its many derivatives—Kubuntu, Xubuntu, Lubuntu, etc.

The Official Full Circle Style Guide can be read at: <http://bit.ly/fcmwriting>

Please read this document before submitting an article. Follow the guidelines and you will have a much better chance of seeing your article in Full Circle.

Writing

There is no word limit for articles, but be advised that long articles may be split across several issues. In your article, please indicate where you would like a particular image to be. Please do not use any formatting in your document.

Images

Images should be no wider than 800 pixels, in JPG format, and use low compression. When you are ready to submit your article, please email it to: articles@fullcirclemagazine.org

Non-English Writers

If your native language is not English, don't worry. Write your article, and one of the proofreaders will read it for you and correct any grammatical or spelling errors. Not only are you helping the magazine and the community, but we'll help you

with your English!


FOR REVIEWS:

Games/Applications When reviewing games/applications, please state clearly:

- *title of the game*
- *who makes the game*
- *is it free, or a paid download?*
- *where to get it from (give download/homepage URL)*
- *is it Linux native, or did you use Wine?*
- *your marks out of five*
- *a summary – with positive and negative points*

Hardware When reviewing hardware, please state clearly:

- *make and model of the hardware*
- *what category would you put this hardware into?*
- *any glitches that you had while using the hardware?*
- *easy to get the hardware working in Linux?*
- *did you have to use Windows drivers?*
- *marks out of five*
- *a summary – with positive and negative points*



[RETURN TO CONTENTS](#)

Review - mtPaint Graphics Editor

Written by Adam Hunt


mtPaint is a simple graphics editor for the Linux desktop and also for Microsoft Windows. It is intended for use in creating pixel art and other simple graphics, as well as for photo editing.

History

mtPaint was started by British programmer Mark Tyler, as a personal project in 2004. He wanted a graphics editor for his own use, and then decided to share it as a free software program, licensed under the GNU General Public License. He started the project on 7 August, 2004, and made the first public release just 37 days later, on 13 September, 2004. Tyler attributed his rapid development to having studied mhWaveEdit, which was written by Magnus Hjorth.

Tyler's goals for mtPaint included that it be fast, simple and lightweight, so it could be run on older hardware. He wrote it in "C" and using GTK1-2, an interface it retains today. That gives it a bit of a "dated" look, but it does have low resource requirements.

Following the first release, version 0.23, Tyler quickly added more features including Windows support following version 0.30. Layers were included in version 2.00. Version 3.00 brought channels, including coding contributions by Dmitry Groshev who wrote the coding for the alpha (transparency), selection and mask channels.



It was at this point, in 2008, that Tyler ended his work on mtPaint to pursue other software projects and Groshev became the lead developer, a role which he continues today.

The most recent release of mtPaint is version 3.40, which dates to 30 December, 2011. It isn't "abandonware", it just seems to have reached a point where no new features are needed, at least for now. At some point, it will probably need the interface rewritten using GTK3 or Qt to modernize it for ongoing use. Meanwhile, it still works well.

mtPaint has been included as a default application in a number of Linux distributions like Puppy Linux. It was also included in Lubuntu until Lubuntu 18.10 – with the move away from GTK applications to Qt-based applications instead. mtPaint can still be installed on Lubuntu or any other Ubuntu flavor.

Features

While being simple and light, mtPaint includes all the tools that most users will need for creating drawings and editing photos. mtPaint supports a number of formats for both import and export: BMP, GIF, JPEG, LSS, PNG, TGA, TIFF, XPM and XBM.

It supports cutting and pasting, scaling and resizing, rotation including arbitrary rotation, inverting colors, grayscale and a range of filters, like sharpen, unsharp mask, soften, Gaussian

blur, emboss and “bacteria”. Images open full-sized, but can be zoomed from 10% to 2000% for working on them. It has 100 possible layers and 81 included brush shapes, although custom shapes can be defined, too. It comes with a large collection of built-in keyboard shortcuts, too, which makes quick work of repeated editing actions.

It is not easy to figure out how to use mtPaint on your own (it has low “discoverability”). It has some odd quirks that are different from other image editors. For instance, when you cut and paste one image onto another, mtPaint won’t “stick” (commit it) until you hit “return” or right-click.

Cropping an image is also unconventional. Instead of marqueeing the portion and then copying and pasting it into a new window, on mtPaint you marquee the portion you want and then hit “delete”, the rest disappears and the cropped image is ready to be saved. Easy, when you know how!

To help new users learn their application, mtPaint has a pretty good user manual, which Groshev recommends reading. He wrote that mtPaint’s features *“may look opaque to users who do not like to read docs; but not everything in image processing can be made self-explanatory, particularly if one tries to keep the program small. Those features which aren’t obvious, have explanations in the handbook; time spent looking them up will be well rewarded by not wasting time on learning things by trial and error.”* Some users have also written on-line tutorials, all of which are useful.

Drawbacks

During my testing of mtPaint 3.40, I found only two things I thought need fixing. One is that images cannot be opened from my file browser in mtPaint. The application opens, but indicates that the image cannot be opened. Opening photos from mtPaint File -> Open works fine, though.

The other issue is that there is no way to preview arbitrary image rotation angles. You have to either guess and try it and then “undo” if it isn’t right, or use the rectangular marqueeing tool to measure the angle from the vertical or horizontal, which

it displays, and then rotate to that angle. That procedure is described in the user manual, too.

Why mtPaint?

Why use mtPaint instead of other applications like GIMP? The question is an obvious one to ask, but mtPaint has some advantages over the more fully-featured GIMP. mtPaint is much lighter-weight, opens much faster, and is simpler to use. It also offers some features that have recently been removed from GIMP like the plain, old “sharpen” filter, which is still useful.

For many users who just want to do some basic photo editing, like scaling the image down, cropping, fixing contrast and sharpening, mtPaint can do the job quickly and without eating up much RAM in the process.

Conclusions

mtPaint is a very mature application – with 15 years of development behind it. It is fast to use, and has enough features that many users will find it very useful for photo editing and making drawings.

Adam Hunt started using Ubuntu in 2007 and has used Lubuntu since 2010. He lives in Ottawa, Ontario, Canada in a house with no Windows.

[RETURN TO CONTENTS](#)

Interview: Jaron Parsons - Fury BSD Dev

Written by SJ Webb

How many members are on the FuryBSD Team?

Right now Joe Maloney is our Primary Developer. The rest of us perform minor development roles; there are currently 3-4 of us, and a growing community willing to help.

What is your name? And your role on the team?

Jaron Parsons.

Primarily, debugging, web/community services management, assist in creative brainstorming on features and improvements, and minor development assistance thus far. As time permits I hope to help more with as much as I can.

What is your background? How did you get started in this field? Any suggestions to others who want to follow your path?

Currently I am Managing IT for iXsystems, Inc. Joe is Managing the Quality Engineering dept at iX as well. Joe and I both have worked together for many years. Joe and I worked for a small ISP in KS where he was the primary support technician and I was managing the dept and handling system administration. We both managed many FreeBSD systems during that time, and ran most of the ISP systems on FreeBSD, and FreeNAS servers. Prior to that, I have been in various IT positions since the mid 90s, and have been a FreeBSD user since the mid to late 90s as well.

What lead you to BSD? Why not another operating system

platform?

My initial draw to BSD was simple curiosity. I had a friend talking about it, and the difference between it and linux. I think that was around the time linux had just started really, and the BSD lawsuit was ending. So I went out and tried to find a copy of the BSD and linux CDs to check them out. Unfortunately, I couldn't find BSD on any shelves, at least not in KS, and this was when the internet was served on a 2400 baud modem. Ha! I did find a copy of Linux slackware at a Hasting bookstore. Might have been the first release, I am not sure now. I installed it, used it, and just loved the feel of the unix-like shell. I finally got my hands on a copy of FreeBSD, and felt it was better organized, and better to use. I continued using it lightly for various things. My first network Firewall Router, first email and web server, and etc. When the ISP hired me, I rebuilt their services from the ground up using FreeBSD for nearly everything. The rest is history. FreeBSD has always been reliable, stable, and fast. In recent years, FreeBSD has had some struggles, just things I have picked up in forums and mailing lists over the years. BSD has been an important part of my career and wish to see it succeed. A FreeBSD desktop has always been something I had hoped for. Specifically one that left the underlying OS as much intact as possible, and built upon existing tools. My hope is FuryBSD will follow this direction. Regardless if it does, I think it will be a fun project.

Who do you see as a common user of BSD?

Currently, mostly Sysadmins or IT professionals. FuryBSD can bring more interest from standard desktop users.

What are the best resources for a new BSD user?

The Internet, Google, many various BSD related communities. The FreeBSD Handbook is your friend as

well ;)

Editing configuration files is daunting for new users. Do you see this issue preventing new user adoption? Any suggestions or sources to help people develop this skill?

For users new to the *NIX world, it can be overwhelming. I think many find interest in the challenge. But, in order to get new Desktop oriented users, it would be best to build tools that make difficult tasks easier. Joe has been working on many tools to help with this, and the feedback and support from the community helps to develop these very much. Currently, it is just too many tasks for the few of us. Hopefully, more will come forward and offer to help. We are just a few guys with a passion for BSD and want to make something of it.

How difficult was it to develop the platform? Why FreeBSD as the base?

Most of the troubles are things that just need attention in FreeBSD base. Why FreeBSD.. I think I answered that above ;)

There is no software center, so everything is done by the command-line and terminal?

Currently, yes. The FreeBSD pkg system isn't hard – once you get it down, but it is a boring command-line tool. We have talked about possible graphical tools, and I am sure it will be something that is revisited again in the future.

Where do you see FuryBSD in the next 3 years?

Hopefully on every desktop! Realistically, I am hoping for a good strong following, and active development moving forward steadily. Possibly fixing some issues which can help the upstream FreeBSD, and tools to help the common PC user to jump into experiencing

BSD with ease.

What is the most common use for FuryBSD?

Currently, I can see it becoming a development workstation, or standard desktop for various business uses, as well as home enthusiast.

How can someone help your project?

Join the community! Start with installing it. Explore, and provide feedback to the group. Submit Bugs, or work on the development of requested features. The more the merrier!

Thank you for your time.

Thank you for the opportunity!

[RETURN TO CONTENTS](#)

Letters

Compiled by Ronnie Tucker

letters@fullcirclemagazine.org




DistroWatch.com

Put the fun back into computing. Use Linux, BSD.

[RETURN TO CONTENTS](#)

Q&A

Compiled by Erik The Unready



Planned obsolescence is something I see with Apple products all the time. Perfectly good hardware, that could last another 20 years, is just pushed aside. Linux is then hailed as the saviour, but modern Linux also does not run on old hardware. I also find that some things are made and just pushed out there and never updated. What is the answer, I wonder? Do we need regulation on software and hardware?:


- *Example 1: One of my clients has an expensive microscope in a lab that works like a dream, connected to a Windows 98 computer. Some idiot stuck their phone into the USB port and gave it a virus. The company that sold it to them, no longer supports the software and they now need to buy a new microscope at tremendous cost to get the new software, that (gasp!) runs on windows 7 (also becoming obsolete!). Since the software was never given to the client, but a technician used to come out from the supplier to reload, they cannot even reload the software themselves.*
- *Example 2: A machine was built in Germany and runs Damn Small Linux. Software, manuals etcetera, was all supplied to the client, but it turns out it was the wrong version (probably upgraded some time thereafter), and a technician had to be sent from Germany to reload it. When the technician arrived*

and tried to install from supplier CDs, he found that version issues prevented him from running the database. Since the machine is offline, and Damn Small Linux is gone, the poor guy had to sit all night on the telephone with a developer to get it to work. If the hardware fails again, this also will be history as DSL does not work on UEFI. Also the hack to make it work on an older version of DSL is an issue. Instead of making a new computer with say, Ubuntu, and selling that to the client, the whole machine has to be replaced at astronomical cost.

Is Linux the answer or does corporate greed just steam-roll everything? Or how does one plan for the future? Does perfectly good equipment need to go to the landfill just because it is tied to software?

Q: Please give assistance with Ninja-ide. I am installing on Ubuntu 18.04. My laptop Lenovo Yoga. Cannot run. Starting in terminal is not finding python2.7. Ninja-ide is best ide. Everyone advises to use Pycharm. I am poor.

A: Your only alternative then, is to build from source. Create a folder and open a terminal in it. Type the following commands:



Now create a launcher for it, as it will not show up in your menu, or just run it from the command-line.

Q: OK, my question is this: I downloaded and installed a DEB file that wasn't in the software centre. How do I uninstall it?

A: The easiest way is to run that .deb-file with Gdebi, and since it will detect that it is already installed, just click the remove button. Otherwise, search in the

software centre, it is probably there and you are just looking right over it. (It should add to the software centre once installed). Of course you can remove it via the terminal too, with apt remove.

Q: I really don't like the terminal or man pages. Whenever I have a question, people say have you read the man pages? This cheeses me off to no avail, makes me want to scream. Why can't I just work in my GUI?

A: Work in your GUI. Nobody is making you do anything. You can read the man pages on the internet in the comfort of your browser. Know there are also alternatives to man pages, like yelp, cheat and bro pages. TL;DR is another alternative. You can also keep a copy of the Linux Bible on your desk.

Q: My Xubuntu is stuck on the screensaver and won't let me in. I turn off my laptop & when I turn it on again, it is still there, stuck. How do I get around it?

A: That is an issue with the screensaver in Xubuntu. But I am going to assume your bigger issue is not being able to kill it. I suspect your fancy laptop has an SSD and turning it off, just hibernates the computer as the drive is too fast. Remove the power and the battery, reinsert the battery and restart and change or disable the screen saver.

Q: I am trying to install Opera on Ubuntu 18.04.2. When I run `sudo apt install opera`, I get dependency not satisfiable. How can I fix this?

A: Grab the opera .deb-file from www.opera.com and install it with gdebi. You may see a message about 'another version being available in the repo's', but ignore that and install it manually. It will ask you if you would like to update opera with the rest of your system, say yes.

Q: You da man, help me get a comic book reader on Ubuntu

18.04. Software centre only has crap when I search for comix...

A: I be da small furry creature from Alpha Centauri, what do you mean? Try: `sudo apt install mcomix` from a terminal.

Q: I'm trying to figure out why I can't install `spez browser` from `deb spezbrowser_9.0_amd64.deb`. Could it be that this was made for Debian / Elementary and not Ubuntu? My hardware is i3 with 8GB of memory.

A: I would venture that the version is too old. Try `Spez browser` version 10. <https://sourceforge.net/projects/spez-browser-mirrors/>

Q: I did something dumb and I don't know how to fix it. I am new to Linux and not comfortable with the command-line. I am running on Voyager, Ubuntu edition, and I used the fancy panel to turn off my trackpad, as it kept interfering with my typing. How can I turn it on again as I set the panel to auto hide and I can't bring it back up.

A: Plug in a mouse and reverse what you did. No command-line needed.

Q: My `dmesg` log is filled with entries like this: `audit: type=1400 audit(1576129910.654:67): apparmor="ALLOWED" operation="open" profile="libreoffice-soffice" name="/home/perry/.thunderbird/cv58948v.default/key3.db" pid=11348 comm="soffice.bin" requested_mask="wr" denied_mask="wr" fsuid=1000 ouid=1000`
What does that mean? Am I hacked from LibreOffice?

A: Nothing to be alarmed about. It is simply a bug. See: <https://bugs.launchpad.net/ubuntu/+source/apparmor/+bug/1849680>

Q: Which version of Ubuntu do you recommend I replace Mac OSX 10.5 with on my G5 mac? I can't seem to install any version. I have verified my iso images.

A: You will need the PowerPC version of an LTS release. Look at PPC in the ISO name of the install CDs.
<https://help.ubuntu.com/lts/installation-guide/powerpc/ch02s01.html>

Q: My laptop is feeling slow with Ubuntu. I am having 8GB of memory, and 500GB of RAM. I have made swap file 16Gb, but no increase performance. <image> ???????? should this say 500 GB memory and 8 GB RAM ?????

A: You don't give me a lot to go on. Increasing the swap file willy-nilly for speed is analogous to downloading RAM. I would say look at readahead, vm swappiness, startup applications, and don't mix environments, like KDE and Gnome. BIOS setup may also be an issue, make sure you are on AHCI and not compatible. There are just too many vectors and too little information provided.

Q: My friend, on Xubuntu, how can I make this?
<screenshot170518.jpg> See thru bar.

A: Right-click on your panel, choose properties, then set the style to transparent.

Q: I am done with Windows 10 \$#!^!. Updates when I want to shut down, junk in my start bar, more junk every day I didn't ask for. Restarting without my permission. Every third day Office365 decides it isn't registered. I was dual-booting, but now I want to stick with Kubuntu only. How can I get rid of this \$ #!^?

A: Well, you can simply edit GRUB to boot only to Kubuntu. If you want to reclaim the space, you can, with gparted. However, I would say, save yourself the trouble and copy your stuff off and reinstall a nice fresh copy of Kubuntu that overwrites Windows. I have seen weird stuff happen after updates when Windows has just been silenced.

Q: Firstly, I like LibreOffice. The problem is my wife does not.

She has only known Microsoft office her whole life. I have heard there is a skin pack that you can make LibreOffice look like MS office. How do I go about doing that?

A: I have never done that. I usually try OnlyOffice or WPS office for my clients who do not like LibreOffice. That said, there are videos on the subject: <https://www.youtube.com/watch?v=Fx-h3n0RWso> and https://www.youtube.com/watch?v=Q2uE_q1YLp8

Q: When I close my lid on ubuntu and open it again hours later, I want to set the brightness, but the fn + F9 key combination does not work. Why?

A: I am going to assume that you mean on the login screen. That is because you have not authenticated as a user of the system and therefore cannot set any system settings. Only once you log in, will this function become available.

Q: I am new to Linux in general. What reading do you recommend?

A: Man, I have heard about this awesome FREE magazine – www.fullcirclemagazine.org, but it may just be rumors.

Q: When I am changing icon themes, my battery icon in my taskbar lies down. How come it doesn't stay vertical?

A: Those icons are theme dependent. The thing with open source is if you do not like something change it. Look in `~/.local/share/icons` . Change them as you please.

Q: First time on Ubuntu studio. My question is this: How do I connect my Xenyx mixer to Ardour?

A: You can slap me with a salmon and rub jam in my armpits if I know. (When it comes to expensive hardware, I am not able to assist. Whatever you paid

for it, times it by 16, and see if you would have paid that for it. I am in that boat). My best guess would be JACK. Read here: <https://help.ubuntu.com/community/UbuntuStudio/UbuntuStudioControls> (maybe a reader has more information here?).

Q: I have a Genius wireless mouse connected to my Lenovo laptop running Ubuntu 18.04. Whenever I write a disk image with balena etcher, the mouse starts to jitter and goes completely mental. Rebooting the laptop helps only sometimes. Turning the system off and coming back later solves it. I think I need to update my BIOS, but I am afraid of bricking it as it is out of warranty.

A: I do not have a Genius mouse to test, but! When writing a large file to USB, it takes a lot of power and the USB is likely to get hot. Heat is the enemy of computer parts. The issue with USB ports is that they are on the same bus, and they all get hot simultaneously. Take out your mouse transceiver and blow on the metal for 30 seconds. Put it back in and you should be good. (I know it sounds hocus-pocus, but it works).

Q: Hey man, something strange started happening. I have a WD 500GB external drive. It is formatted with NTFS. All of a sudden, I can't write to it. I see no errors in dmesg. Google searches have me editing fstab, which I never needed to. Even as root, I cannot write to the drive. Gnome 'disks check' says nothing is wrong. Can you help me? I am still on Ubuntu 18.04. Nothing has changed and all the permissions are still intact.

A: Terminal time! (doesn't quite sound the same as hammer time): Open your terminal and type: `sudo parted -l` (to get the drive name)
Now unmount that drive and type: `sudo ntfsfix /dev/sdXY` (XY being YOUR drive numbers where your NTFS partition is, for instance: sdb1). Mount your drive normally. It *should be sorted out.

Q: My problem is with VLC on ubuntu. My laptop screen is 1366x768 and when I use VLC skins, it becomes bigger than my screen and I can't resize it. <electrix.vlt>

A: As a fellow WXGA screen user, I feel your pain. My only advice to you is don't use that skin. (I am not being nasty here, but those skins are ten years old and VLC has changed a lot over the years). You can make your own. <https://www.videolan.org/vlc/skineditor.html>
<https://www.videolan.org/vlc/skineditor.html>

Q: bonjour, antivirus clam tk a détecter 57 possibilité viral dans votre suite bureautique libre office, je ne comprend pas pourquoi cela c'est possible étant donné que libreoffice est intégré de base a votre système, je n'ai pas supprimer les infection signalée par l'antivirus j'ai des doute et que ci je supprime tout cela que cela pose problème a la suite bureautique a par cela Ubuntu 18.04 LTS 32-bit fonctionne très bien en vous remerciant pour cette version moi qui possède encore un ancien ordinateur car je n'ai pas encore les moyen dans avoir un de nouvelle génération. Je souhaite a toute l'équipe de UBUNTU des bonnes fête de fin d'année

A: Nous sommes un groupe de bénévoles et non Ubuntu. Vous pouvez simplement supprimer le redémarrage de LibreOffice et le réinstaller. Je doute qu'il y ait un virus, mais ne prenez aucun risque. Lisez notre article sur Lynis et utilisez-le pour renforcer votre ordinateur.

Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.

Linux Certified - Part 9

Written by Erik the Unready

Maybe next month

[RETURN TO CONTENTS](#)

Ubuntu Games - Stygian – Reign of the old ones

Reviewed by Erik

Website: https://www.gog.com/game/stygian_reign_of_the_old_ones

Price: \$29

When I say H.P. Lovecraft, you immediately think of Cthulhu? Right? What if I told you Lovecraft mentioned Cthulhu only in Mountains of Madness? Almost everything else you read about the Cthulhu mythos is a community project? This is a shining example of what the community can do. Countless fan fiction that adds to the mythos. Which is another reason I like the Cthulhu mythos, there are so many possibilities. Building on this mythos, is a game at GOG that is Linux compatible. It can be yours for the low-low price of \$30, ouch! What?



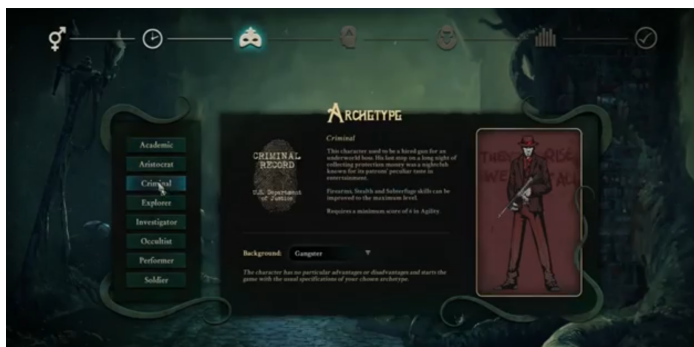
Let us see if this game lives up to the price tag.

Story:

The adventure is set in the town of Arkham. I'm the Batman! - No, not that Arkham. Here the town is a 1920's-type town (or shall I say Victorian era?) that is also trapped in another dimension. As with most of the Cthulhu mythos, this story revolves around your sanity. So, it is no wonder the story starts with you in bed, and a creepy character watches you sleep. You follow this character into the shadier parts of town. Until the 'shrooms kick in and you start seeing things. Okay now....

Gameplay:

Like with most role-play games, you need a character. You can choose from eight pre-made ones, or you can roll your own. The nice thing about this character creation mechanism is that you get to choose things like age, and it has an effect. Backgrounds are not just for show either, they too have an effect. Should you choose the pre-made Sherlock Holmes character, you get proficiencies in guns, investigation, stealth and psychology. This affects your speech options with NPC's as well, not just combat. Speaking of combat, you do not have to fight immediately, you can sneak and loot and barter. Character perception is investigation and the higher the 'stats', the easier you will find hidden items, which can be a real boost. It kind of reminds me a bit of Fallout 1 & 2, with different graphics, though there is no fast travel. The Archetypes screen offers you a choice of Academic, Aristocrat, Criminal, Explorer, Investigator, Occultist, Performer and Soldier. This will affect how NPC's respond to you or barter with you. It will also affect how your sanity is measured. The combat is hex based and turn based, with interesting animations. Your actions are determined in action points. You obviously have weapons, but you are armed with spells as well. Different spells, however, have different penalties, for instance blood magic may cost you health as well. You will find obstacles on the "battlefield" too, but they do not provide cover, instead they hinder movement. Sanity, sanity, sanity. You will lose sanity like clockwork. Then, to regain some sanity, you will have to trade off something else, like health. This game is all about choices and lets you think about your choices.




Graphics and sound:

The art style seems hand drawn (maybe?), and the world seems hand-crafted. The only thing that had no consistency was the portraits, though they gave a 1920's feel. The style had me thinking that the game was a visual novel in the beginning and retains the comic book feel throughout. The voice acting is on the money. There just seems to be a shortage of it. The music and atmospheric sounds fit the game like a glove. Very immersive, to say the least.

Conclusion:

There are some bugs, as the game is quite new. The forums have things like giving a party member something and not being able to get it back. That said, the story seems great (I say seems as I have not finished it yet), the dialogue seems to drive the story, and not every outcome is favourable. Again this makes you think about your choices, unlike other games where you can go through all the choices. The way the dialogue is structured, I suspect there are multiple endings. Character creation options are plentiful and once you have played, you will realise why. You may even be stuck there for a while. 1C is known for their "war" games, and this makes a welcome change. \$30 in local monies is a bit much, so I will wait for a sale.

This game is a must for role play / puzzle / Cthulhu fans. I am Cthulhu'd up and will look at Gibbous next.



Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.

[RETURN TO CONTENTS](#)



The current site was created thanks to **Lucas Westermann** (Mr. Command & Conquer) who took on the task of completely rebuilding the site, and scripts, from scratch, in his own time. The Patreon page is to help pay the domain and hosting fees. The yearly target was quickly reached thanks to those listed on this page. The money also helps with the new mailing list that I set up.

Several people have asked for a PayPal (single donation) option, so I've added a button to the right side of the website.

A big thank you to all those who've used Patreon and the PayPal button. It's a HUGE help.




<https://www.patreon.com/fullcirclemagazine>



<https://paypal.me/ronnietucker>




<https://donorbox.org/recurring-monthly-donation>



Patrons and Donors

Monthly Patrons





Single Donations

2019:



2020:





[RETURN TO CONTENTS](#)

How to Contribute

FULL CIRCLE NEEDS YOU!

*A magazine isn't a magazine without articles and Full Circle is no exception. We need your opinions, desktops, stories, how-to's, reviews, and anything else you want to tell your fellow *buntu users.*

Send your articles to:

articles@fullcirclemagazine.org

*We are always looking for new articles to include in Full Circle. For help and advice please see the **Official Full Circle Style Guide** : <http://bit.ly/fcmwriting>*



FCM #154 Deadline:
Sunday 9th February
2020. Release: Friday
28th February 2020.

- Send your comments or Linux experiences to:
letters@fullcirclemagazine.org
- Hardware/software reviews should be sent to:
reviews@fullcirclemagazine.org
- Questions for Q&A should go to:
questions@fullcirclemagazine.org Desktop screens should be
emailed to: misc@fullcirclemagazine.org
- ... or you can visit our site via: fullcirclemagazine.org

ronnie@fullcirclemagazine.org

admin@fullcirclemagazine.org

FCM PATREON : <https://www.patreon.com/fullcirclemagazine>

Getting Full Circle Magazine:



EPUB Format - Most editions of Full Circle have a link to the epub file on the downloads page. If you have any problems with the epub file, you can drop an email to:

mobile@fullcirclemagazine.org



Magzster - You can also read Full Circle online via Magzster: <http://www.magzster.com/publishers/Full-Circle>. Please share and rate FCM as it helps to spread the word about FCM and Ubuntu Linux.



Issuu - You can read Full Circle online via Issuu: <http://issuu.com/fullcirclemagazine>. Please share and rate FCM as it

helps to spread the word about
FCM and Ubuntu Linux.

For the Full Circle Weekly News:



You can keep up to date with
the Weekly News using the RSS
feed: [http://
fullcirclemagazine.org/feed/
podcast](http://fullcirclemagazine.org/feed/podcast)



Or, if you're out and about, you
can get the Weekly News via
Stitcher Radio (Android/iOS/
web): [http://www.stitcher.com/
s?fid=85347&refid=stpr](http://www.stitcher.com/s?fid=85347&refid=stpr)
and via TuneIn at: [http://
tunein.com/radio/Full-Circle-
Weekly-News-p855064/](http://tunein.com/radio/Full-Circle-Weekly-News-p855064/)



[RETURN TO CONTENTS](#)